

The Value of Intrusion Detection Systems in Information Technology Security Architecture

Huseyin Cavusoglu

A. B. Freeman School of Business, Tulane University, 7 McAlister Drive, Goldring/Woldenberg Hall,
New Orleans, Louisiana 70118, huseyin@tulane.edu

Birendra Mishra

School of Management, University of Texas at Dallas, Richardson, Texas 75083, and
Anderson Graduate School of Management, University of California, Riverside, Riverside, California 92521,
barry.mishra@ucr.edu

Srinivasan Raghunathan

School of Management, University of Texas at Dallas, Richardson, Texas 75083, sraghu@utdallas.edu

The increasing significance of information technology (IT) security to firms is evident from their growing IT security budgets. Firms rely on security technologies such as firewalls and intrusion detection systems (IDSs) to manage IT security risks. Although the literature on the technical aspects of IT security is proliferating, a debate exists in the IT security community about the value of these technologies. In this paper, we seek to assess the value of IDSs in a firm's IT security architecture. We find that the IDS configuration, represented by detection (true positive) and false alarm (false positive) rates, determines whether a firm realizes a positive or negative value from the IDS. Specifically, we show that a firm realizes a positive value from an IDS only when the detection rate is higher than a critical value, which is determined by the hacker's benefit and cost parameters. When the firm realizes a positive (negative) value, the IDS deters (sustains) hackers. However, irrespective of whether the firm realizes a positive or negative value from the IDS, the IDS enables the firm to better target its investigation of users, while keeping the detection rate the same. Our results suggest that the positive value of an IDS results not from improved detection per se, but from an increased deterrence enabled by improved detection. Finally, we show that the firm realizes a strictly nonnegative value if the firm configures the IDS optimally based on the hacking environment.

Key words: economics of IT security; intrusion detection systems (IDSs); ROC curves; security configuration; IT security management

History: Tridas Mukhopadhyay, Senior Editor; M. S. Krishnan, Associate Editor. This paper was received on December 5, 2001, and was with the authors 5 months for 2 revisions.

1. Introduction

Dramatic increases in the number of IT security breaches and resulting monetary losses in recent years have made IT security a top issue in the management of IT infrastructure,¹ which is also reflected in

¹ The number of computer intrusion cases filed with the Department of Justice jumped from 547 in 1998 to 1,154 in 1999 (Goodman and Brenner 2002). The losses from computer crime incidents reported by the Computer Security Institute (CSI)/Federal Bureau of Investigation (FBI) surveys were \$456 million in 2002, in contrast to \$378 million in 2000 and \$266 million in 1999 (Power 2002). A global survey conducted by InformationWeek and Pricewaterhouse Coopers LLP estimated that computer viruses and hacking took a \$1.6 trillion toll on the worldwide economy and a \$266 billion toll in the United States alone (Denning 2000).

the increasing security budgets of firms (Hulme 2002). Businesses and governments have undertaken several measures to minimize the loss from security breaches. IT security-related laws, popularly known as cyber laws, enacted by governments, act as broad deterrents against IT-related crimes. These external control mechanisms supplement a firm's internal control mechanisms. Traditionally, internal controls fall into two major categories: preventive and detective. In the IT security context, preventive controls, such as firewalls, aim to develop a defensive shield around IT systems to secure them from intrusions. Detective controls, such as IDSs, try to detect intrusions that have already occurred. Because complete pre-

vention of intrusions is unlikely, detective controls have become an important element in a firm's overall security architecture.² Furthermore, studies have reported that many hackers are employees or insiders (Escamilla 1998, Russell and Gangemi 1992). Detective controls complement preventive controls by identifying intrusions from both insiders and outsiders.

Intrusions are caused by external hackers accessing the system using the Internet, or authorized external and internal users attempting to gain additional privileges or to misuse their privileges. "Intrusion detection systems (IDSs) are hardware or software systems that automate the process of monitoring the events occurring in a computer system or network, analyzing them for signs of security problems" (Bace and Mell 2001). These systems warn security experts about suspected intrusions. However, IDSs are imperfect; they have false positive errors (warning when no intrusion exists) and false negative errors (no warning when an intrusion exists). Consequently, security experts manually investigate log files and audit trails of users that generate alarms from an IDS to confirm or rule out intrusions. Security experts may also investigate logs that do not generate alarms from the IDS to detect missed intrusions.

Despite the economic importance of IT security to organizations, academic research has not analyzed the value of IT security technologies. Rather, the bulk of academic security research has focused on developing algorithms to be implemented in security technologies, and improving their efficiency and effectiveness.³ Although the value of preventive controls, such as firewalls, is obvious, organizations are still uncertain about the value of detective controls, such as IDSs. Among practitioners, the value of an IDS to firms that deploy it has generated significant attention (Shiple 1999). The proponents of IDSs claim

that because technology that prevents intrusions does not exist, IDSs may be the only efficient recourse to deal with intrusions that have already occurred. The opponents of IDSs claim that IDSs suffer from low detection and high false alarm rates. Given that the proportion of hackers in the user population is generally low, an IDS with even a moderate false alarm rate generates more alarms for normal users than for hackers. This phenomenon, known as the base-rate fallacy (Axellson 2000), often causes firms to ignore alarms from an IDS, which may render the IDS useless. A recent report from Gartner Inc. (Gartner 2003) intensified the debate when it dismissed IDSs as failed technology and recommended that firms allocate all their security budgets to preventive controls. To eliminate the uncertainty surrounding the value of IDSs and to guide firms in their IDS implementations, a rigorous assessment of the costs and benefits of IDSs is critical.

Another issue confronting a firm that deploys an IDS relates to its configuration. The quality profile of an IDS is measured by its false positive and false negative rates. Although having both rates low in an IDS is desirable, the technology is such that a reduction in one type of error is often accompanied by an increase in the other type. The goal of configuration is to balance the two error rates to minimize the firm's cost. Guidelines from commercial security software manufacturers⁴ and research institutes, such as the Software Engineering Institute (SEI) at Carnegie Mellon University, emphasize the need for proper configuration of IDSs. For example, SEI's report on IDSs (Allen et al. 2000) cautions firms against accepting the default settings automatically and recommends appropriate configuration to balance security and operational requirements. It is imperative that a firm understands the effect of IDS configuration on the value derived from the IDS to determine the optimal configuration for its operating environment.

In this paper, we seek to offer insights into the two above-mentioned issues: (i) whether and under what conditions an IDS offers value, and (ii) the effect of configuration on the value derived from the IDS. We

² Intrusion detection systems have been the fastest growing security product in terms of sales in recent years (Alpert 1999). According to International Data Corporation (IDC), the market for IDSs has grown from about \$20 million in 1997 to \$100 million in 1999 and is projected to reach \$528 million by 2005 (Messmer 1999). Axent Technologies and Internet Security Systems are the market leaders in the IDS market. The other major players are Cisco, Computer Associates, IBM, and Network Associates.

³ In the computer science literature, Lee et al. (2002) propose building cost-sensitive intrusion detection models using the decision theory approach.

⁴ For example, Sriram (2002) discusses how to choose a threshold value to detect attacks by computer viruses in Novell Border-Manager.

believe that ours is the first study that investigates the value of an IT security technology from an economic perspective.

IT security has long been viewed as a game between firms and hackers.⁵ Thus, we model the IT security problem as a game between a firm that attempts to minimize loss from a security breach and a hacker that wishes to compromise the firm's information system. In our game, the firm decides whether or not to implement an IDS based on its value. In both IDS and no-IDS cases, the firm chooses its manual investigation strategy, and the hacker chooses his or her hacking strategy. When the firm decides to implement an IDS, we consider two scenarios. In the first scenario, which we refer to as the *out-of-box configuration*, the firm does not optimize the configuration. In the second scenario, the firm uses the optimally configured IDS. We summarize the significant findings of our analysis as follows.

(i) In the out-of-box configuration, the firm may realize a positive or negative value from an IDS. The firm realizes a positive value if and only if the detection rate (i.e., true positive rate) of the IDS is greater than a critical value, which is determined by the user's utility parameters.

(ii) In the out-of-box configuration, the firm realizes a positive value if and only if the IDS deters hackers, that is, the hacking probability is lower with than without the IDS. Irrespective of the value, however, an IDS reduces the effective manual investigation rate, thus reducing the manual investigation cost, but does not change the effective detection rate. These results suggest that the value of an IDS arises from deterrence rather than improved detection.

(iii) In the optimally configured IDS, the firm realizes a strictly nonnegative value. The optimal detection rate depends not on the firm's internal cost parameters, but on the external user parameters. In addition, the optimal configuration always deters hackers.

⁵ The game-theoretic aspect of IT security was first noted by Jajodia and Miller (1993, p. 85), "Computer security is a kind of game between two parties, the designer of a secure system, and a potential attacker." Bashir et al. (2001, p. 30) refer to this as the "cat-and-mouse game" between the hacker and the firm. An excellent demonstration of how firms and hackers play the game can also be found at http://www.msnbc.com/modules/hack_attack/hach.swf.

1.1. Related Work

Though we are unaware of a similar study in the IT literature, researchers in other areas have investigated related problems. Our work is most closely related to a broad area in game theory literature known as inspection games. An inspection game is a mathematical model of a situation where an inspector verifies that another party, called an inspectee, adheres to certain legal rules. Researchers have investigated inspection games in areas such as arms control and disarmament, accounting and auditing, environmental control, and crime and punishment.

In arms control and disarmament games, the focus was on detection and verification of treaty violations such as the Non-Proliferation Treaty and disarmament treaty for nuclear weapons (Maschler 1966, 1967; Kilgour 1992; Weissenberger 1992). In accounting, inspection games have been studied in auditing and insurance contexts to deal with the moral hazard that the inspectee may commit irregularities. The emphasis of the research has been on auditors' decision rules for various audit sampling outcomes (Fellingham and Newman 1985, Newman et al. 1996) and on the design of contracts between principals and agents (Baiman 1982, Kanodia 1985, Dye 1986). Relevant research within the area of environmental control analyzed the games between firms that pollute the environment and environmental protection agencies that monitor the activities of firms (Mishra et al. 1997, Russell 1990). The inspection games considered in the crime control literature include patrol of smuggling activities (Thomas and Nisgav 1976), policing of theft and pilferage activities (Feichtinger 1983), and determination of optimal penalties for effective crime deterrence (Becker 1968, Stigler 1970, Polinsky and Shavell 1979, Sethi 1979, Shavell 1991, Mookherjee and Png 1992).

Our work differs from prior work on inspection games in two respects. First, prior work viewed crime prevention or detection technology as a black box, and consequently did not model the technology. Because we are interested in the value of an IDS and its configuration, we explicitly model the technology. Second, unlike the models used in prior literature, our inspection decisions are not based only on random sampling. Instead, inspection decisions are

based on whether or not the IDS generates alarms. In essence, the IDS performs the first set of inspections and the manual audit performs the follow-up verification. Our modeling of the imperfect nature of the IDS technology, coupled with manual inspections that are the norm in IT security domains, yields new insights that are particularly relevant for the IT security domain.

The rest of this paper is organized as follows. In the next section, the IDS technology is explained in detail to build a background on the security model that is presented in §3. Section 4 derives the primary results of the model. Section 5 derives results concerning the value of IDS. We discuss implications of our results and extensions to our basic model in §6. Section 7 concludes the paper with a discussion of the limitations of our research and future research directions.

2. An Overview of IDS Technology

IDSs are software or hardware systems that monitor the events occurring in computer systems and that warn human experts about suspected intrusions⁶ (Amoroso 1999). An IDS uses audit trails and network packets to detect intrusions. Audit trails store information such as date and time of the event, type of the event, origin of the request, and objects accessed, modified, or deleted (National Computer Security Center 1988). IDSs use two types of analysis to detect attacks: signature-based detection, and anomaly detection (McHugh et al. 2000). Signature-based detection looks for events that match a predefined pattern of events, called signatures, associated with a known attack.⁷ Signature-based detectors are effective in detecting common forms of attacks without generating an overwhelming number of false alarms. A limitation of signature-based detectors is that they can only detect those attacks they know about. They must also be constantly updated with signatures of new attacks. Anomaly detection identifies abnormal behavior. Anomaly detection techniques

use a “normal activity profile” for a system and flag all system states varying from the normal profile in a statistically significant manner. The normal activity profiles are constructed typically from historical data.⁸ Unfortunately, anomaly detection often produces a large number of false alarms because normal patterns of users and system behavior can vary widely. However, unlike signature-based IDSs, anomaly-based IDSs are capable of detecting unseen attacks.

2.1. IDS Quality Profiles and Receiver Operating Characteristics Curves

The quality profile of an IDS, measured by its false positive and false negative rates, depends on the technology used (signature-based versus anomaly-based), design parameters (for example, the acceptable noise level in an anomaly-based system), and the configuration (strict versus loose). Studies have reported that even the best IDSs could only detect about 80% of the attacks (Lippmann et al. 2000b). They also generate a significant number of false alarms (Lippmann et al. 2000a, b). The quality profile of an IDS is best illustrated and modeled using statistical decision theory. In the simplest case, there are two types of sources that generate inputs to an IDS: normal user and hacker. The goal of the IDS is to classify each user as a normal user or a hacker by examining the user’s transaction history. Two types of errors can occur in this classification: classification of a hacker as a normal user (false negative) and classification of a normal user as a hacker (false positive). We define:

Probability of detection = $P_D = P(\text{classify as a hacker} \mid \text{user is a hacker})$;

Probability of false negative = $1 - P_D$;

Probability of false positive = $P_F = P(\text{classify as a hacker} \mid \text{user is a normal user})$.

In a perfect IDS, P_D will be equal to one, and P_F will be equal to zero. However, as shown in the next two paragraphs, the IDS detection technology is such that a high value of P_D also entails a high value of P_F . This is because of the variability of data associated with normal and abnormal transactions and imprecision of

⁶ IDSs have been an active research area for more than a decade. It started with the seminal paper of Dorothy Denning (1987).

⁷ The algorithms used in signature-based systems are discussed in Garvey and Lunt (1991), Porras and Kemmerer (1992), Ilgun (1992), Lunt (1993), Kumar and Spafford (1996), and Monroe and Rubin (1997).

⁸ The algorithms employed in anomaly-based systems are discussed in Lunt and Jagannathan (1988), Lunt (1990, 1993), Lunt et al. (1992), D’haeseleer et al. (1996), Porras and Neumann (1997), Frincke et al. (1996), Neumann and Porras (1999), and Zamboni and Spafford (1999).

algorithms used by IDSs. Typically, the IDS manufacturer sets a P_D and P_F pair.⁹ These settings are known as the out-of-box or default configuration. In many cases, the deploying firms will be able to change these values, though not independently, through a process called configuration, or tuning (Allen et al. 2000). The quality profile of a configurable IDS, i.e., the possible values of P_D and P_F pairs for the IDS, is characterized by a curve known as the receiver operating characteristics (ROC) curve. The ROC analysis was originally developed in the field of radar signal detection during World War II (Trees 2001). Later, it was adopted by psychology and other research streams (Lippmann et al. 2000a). The ROC curve of an IDS shows the trade-off between the P_D and P_F values of an IDS. In general, configuring an IDS to operate at a higher detection rate, P_D , will result in a higher false alarm rate, P_F , and vice versa.

The ROC curve of an IDS can be derived experimentally or analytically (Durst et al. 1999; Lippmann et al. 2000a). Many IDSs classify a user based on whether a numerical score computed from the transaction history exceeds a threshold value, or whether the transaction data satisfy a set of rules, or both. Consider an IDS that uses a numerical score x computed from transaction data and a threshold value t to detect hackers. Let the IDS classify a user as a hacker if $x > t$ for that user. It follows that

$$P_D = \int_t^\infty f_H(x) dx \quad \text{and} \quad P_F = \int_t^\infty f_N(x) dx,$$


where $f_N(x)$ and $f_H(x)$ are the probability density functions of x for normal users and hackers, respectively. Figure 1 illustrates these probability calculations.

The shape of the ROC curve depends on the probability density functions. We assume that the numerical score used to distinguish normal users and hackers follows an exponential distribution. Exponential distributions, besides being analytically tractable, capture the skewed nature of transaction data very well.¹⁰ If the numerical scores for the normal users and hackers follow exponential distributions with parameters

⁹ The P_D and P_F combination is fixed by the threshold in the analysis engine or by the signature dataset for known attacks, or both.

¹⁰ In §6, we show that our results hold for a wide variety of probability distributions.

Figure 1 Computation of P_D and P_F



θ_N and θ_H , $\theta_N > \theta_H$, respectively, then we can write P_D and P_F as

$$P_D = \int_t^\infty \theta_H e^{-\theta_H x} dx = e^{-\theta_H t} \quad (1)$$


$$P_F = \int_t^\infty \theta_N e^{-\theta_N x} dx = e^{-\theta_N t}. \quad (2)$$

P_D can be expressed as a function of P_F as

$$P_D = P_F^r, \quad (3)$$

where $r = \theta_H/\theta_N$ is between zero and one. Equation (3) represents the ROC curve of an IDS. Figure 2 shows sample ROC curves for various values of r . The parameter r captures the quality of the technology used by the IDS. A lower r results in a steeper ROC curve, which represents a higher quality IDS because for a given P_D (P_F), an IDS with a lower r value has a lower (higher) P_F (P_D). Note that the ROC curve lies in the region above the line connecting (0, 0) and (1, 1)

Figure 2 ROC Curves



in the diagram, which suggests that the IDS performs better than random guessing in detecting intrusions, that is $P_D > P_F$. The diagonal line that connects $(0, 0)$ and $(1, 1)$ will be the ROC curve when the IDS uses random guessing.

An IDS can be configured (or tuned) to find the best operating point within its quality profile, characterized by the ROC curve, that fits the operating risk and cost environment. Configuration can be achieved by fine-tuning attack profiles, in case of signature-based detection, and changing severity levels in the alarms, in case of anomaly-based detection (Internet Security Systems 2001, Panko 2003, Ptacek and Newsham 1998). We describe our model in the next section.

3. Model Description

Our goal is to analyze the value of IDSs using a parsimonious model that captures the essence of a typical IT security environment discussed in the previous section. In §6.2., we relax many of the assumptions of our basic model and show that the results from the basic model are robust. We discuss the three broad components of our model—user, firm, and technology—in the following paragraphs.

User. The user population for our model includes every user of the system being monitored by the IDS. Previous studies have shown that incentives for intruders may be related to a financial gain as well as curiosity, self-esteem, vandalism, peer approval, public attention, and politics (Shaw et al. 1999, Koerner 1999, Rothke 2000). We assume that a user committing the intrusion derives a benefit of μ if the intrusion is undetected. If the intrusion is detected, the user incurs a penalty of β for a net benefit of $(\mu - \beta)$.¹¹ We assume that $(\mu - \beta) \leq 0$; that is, a hacker that is detected does not enjoy a positive utility. The penalty can take different forms, such as legal prosecution or social humiliation. Users choose to hack depending on factors such as the benefit they derive from hacking, the penalty they will receive if they are caught,

¹¹ We assume a homogeneous user population in the paper. We also considered a model that had two types of users, honest and dishonest, where a fraction λ of the users was dishonest and honest users did not hack. Although the analytical results were different, the qualitative results for this model were identical to what we present in this paper.


and the likelihood that they will be caught. We denote the probability that a user hacks by ψ .

Firm. Firms confirm or rule out intrusions through manual investigation and analysis of audit trails, whether or not they have implemented an IDS (NIST 1996, p. 223; McHugh et al. 2000). In general, manual investigation is too costly to be done all the time. When the firm does not deploy an IDS, the security team may manually investigate a proportion (ρ) of users. When the firm deploys an IDS, the security team may investigate a proportion (ρ_1) of users that generate alarms and a proportion (ρ_2) of users that do not generate alarms. The firm incurs a cost of c each time it performs a manual investigation. We assume that manual investigations confirm or rule out intrusions with certainty.¹² When an intrusion is undetected, the firm incurs a damage of d . If the firm detects an intrusion, the firm prevents or recovers a fraction, $\phi \leq 1$, of d . It is reasonable to assume that $c \leq \phi d$, so that the firm's cost of investigation is not higher than the benefit it gets if it detects an intrusion. The damage recovery rate, ϕ , could depend on several factors, including the investment in or effectiveness of manual investigation and damage. As we show in §6.2., dependence among these parameters does not change the essential results of the paper. Most companies estimate possible damages in the risk assessment phase prior to implementing and configuring the IDS (Peltier 2001, Tudor 2001).

IDS. Firms may use an IDS to target users for manual investigation. As discussed in §2, IDSs generate alarms if they suspect intrusions. We model the effectiveness of an IDS through its ROC curve, as discussed in §2.

We model the intrusion detection problem as a game between a firm trying to protect its system and the users trying to breach the firm's system. The objective of the firm is to minimize its expected loss from intrusions. The firm decides whether it should implement an IDS, and, if it should, determines the configuration of the IDS, if possible. The users maximize their expected benefit. The timeline for the game is shown in Figure 3.

¹² We have performed the analysis for the case when manual investigation detects or rules out intrusions only with a probability less than one. Our results do not change qualitatively.

Figure 3 The Timeline for the Game

If the firm does not use an IDS, it relies solely on manual investigations to detect intrusions. When the firm uses an IDS, it uses the IDS and manual investigations. We analyze two scenarios when the firm uses an IDS. The firm may choose to use the default configuration for the IDS, which may be the only option available to the firm if the IDS is nonconfigurable. In the case of a configurable IDS, the firm may configure the IDS prior to deployment.

We assume that the firm and users are risk neutral. We analyze the effect of other risk profiles on our results in §6.2. We provide a summary of notations in Table 1.

4. Model Analysis

We perform our analysis using backward induction. That is, we first derive the equilibrium in the firm's manual investigation and the user's hacking strategies, given that the firm has decided whether or not to implement an IDS. Subsequently, we determine if the firm will implement an IDS, based on the cost in each case.

Table 1 List of Notations

Parameters	
d	Damage caused by an undetected intrusion
c	Cost of manual investigation
μ	Utility of intrusion for users
P_D	Probability of getting an alarm from IDS for an intrusion
P_F	Probability of getting an alarm from IDS for no intrusion
ϕ	Fraction of damage prevented or recovered by the firm when an intrusion is detected
Strategic variables	
ψ	Probability of intrusion by a user
ρ	Probability of manual investigation when there is no IDS
ρ_1	Probability of manual investigation when the IDS generates an alarm
ρ_2	Probability of manual investigation when the IDS does not generate an alarm

4.1. No-IDS Case

In this section, we analyze the case in which the firm does not deploy an IDS. We characterize the game in strategic (normal) form in Table 2. A user's strategy, S^U , is to hack, H , or not hack, NH , i.e., $S^U \in \{H, NH\}$. The firm's strategy, S^F , is to investigate, I , or not investigate, NI , the user, i.e., $S^F \in \{I, NI\}$. In Table 2, the first element in each ordered pair is the firm's cost and the second element is the user's payoff.

We use Nash equilibrium as the solution concept. It is a pair of strategies, denoted as (firm's strategy, user's strategy), such that no player can increase his payoff by unilaterally changing his strategy. At least one solution exists for the game, although this solution may not be in pure strategies. In other words, each player may play a mixed strategy by randomly choosing from his pure strategies according to a probability distribution.

We derive the mixed strategy Nash equilibrium. That is, we solve the game as if the strategy space for the user is $\psi \in [0, 1]$ and the strategy space for the firm is $\rho \in [0, 1]$. The firm's expected cost is

$$F(\rho, \psi) = \rho c + \rho\psi(1 - \phi)d + \psi(1 - \rho)d. \quad (4)$$

A user's expected benefit is

$$H(\rho, \psi) = \psi\mu - \psi\rho\beta. \quad (5)$$

The firm minimizes $F(\rho, \psi)$ and the user maximizes $H(\rho, \psi)$. The solution to this game is stated in the following proposition. (The proofs for the results are available from the authors.)

Table 2 Game for the No-IDS Case in Strategic Form

Firm's strategies	User's strategies	
	H	NH
I	$(c + (1 - \phi)d, \mu - \beta)$	$(c, 0)$
NI	(d, μ)	$(0, 0)$

PROPOSITION 1. *The Nash Equilibrium for the no-IDS case is given by the mixed strategy profile ($\rho = \mu/\beta$, $\psi = c/(d\phi)$).*

The mixed strategy equilibrium given in Proposition 1 shows that the user's optimal strategy of hacking with a probability $c/(d\phi)$ makes the firm indifferent between investigating and not investigating. Similarly, the firm's optimal investigation probability μ/β makes the user indifferent between hacking and not hacking. Mixed strategies can also be given a temporal interpretation. Thus, with a frequency proportional to ψ , the user will hack the system on some randomly selected occasions, and the firm, with frequency proportional to ρ , will inspect on some randomly selected occasions.

It is easy to verify that the probability of manual investigation is increasing in the hacker's benefit from undetected hacking μ , but decreasing in penalty β . Similarly, the probability of hacking is increasing in manual investigation cost c and decreasing in the amount of damage recovered by the detection of the intrusion, ϕd . The firm's expected optimal cost in the case of no IDS is c/ϕ .

4.2. The IDS Case

Table 3 shows the payoffs to the players in the IDS case. The user's strategies remain the same as in the no-IDS case. The strategy space of the firm is more complex because the IDS separates the firm's information into two sets, alarm and no-alarm. The firm has two actions, investigate or not investigate, available in both alarm and no-alarm states. Thus, the strategy space for the firm is the Cartesian product of the actions available at each of these two information sets. That is, $S^F \in \{(I, I), (I, NI), (NI, I), (NI, NI)\}$, where the first element in each pair specifies the firm's action when the firm observes an alarm from the IDS, and

the second element is the firm's action when it does not observe an alarm from the IDS. For example (I, NI) implies that the firm investigates the user if it receives an alarm from the IDS for that user and does not investigate if it does not receive an alarm. Again, the first element in each ordered pair in Table 3 represents the firm's cost, and the second element captures the user's payoff. Though P_D and P_F are related by the equation $P_D = P_F^r$, as discussed in §2, we use both parameters in Table 3 for notational convenience and clarity.

Again, we derive the mixed strategy Nash equilibrium. We solve the game as if the strategy space for the user is $\psi \in [0, 1]$ and the strategy space for the firm is $(\rho_1, \rho_2) \in [0, 1] \times [0, 1]$. The following probability computations are used in deriving the equilibria.

$$\eta_1 = P(\text{intrusion} \mid \text{alarm}) = \frac{P_D \psi}{P_D \psi + P_F (1 - \psi)} \quad (6)$$

$$\eta_2 = P(\text{intrusion} \mid \text{no-alarm}) = \frac{(1 - P_D) \psi}{(1 - P_D) \psi + (1 - P_F) (1 - \psi)} \quad (7)$$

$$P(\text{alarm}) = P_F + \psi(P_D - P_F) \quad (8)$$

$$P(\text{no-alarm}) = 1 - P_F - \psi(P_D - P_F) \quad (9)$$

$$P(\text{hacker is detected}) = \rho_1 P_D + \rho_2 (1 - P_D). \quad (10)$$

The firm's expected cost for the alarm and the no-alarm states respectively are

$$F_A(\rho_1, \psi) = \rho_1 c + \eta_1 (1 - \rho_1) d + \eta_1 \rho_1 (1 - \phi) d \quad (11)$$

$$F_N(\rho_2, \psi) = \rho_2 c + \eta_2 (1 - \rho_2) d + \eta_2 \rho_2 (1 - \phi) d. \quad (12)$$

The firm's overall expected cost is

$$F(\rho_1, \rho_2, \psi) = (P_F + \psi(P_D - P_F)) F_A(\rho_1, \psi) + (1 - P_F - \psi(P_D - P_F)) F_N(\rho_2, \psi). \quad (13)$$

A user's expected payoff is

$$H(\rho_1, \rho_2, \psi) = \psi \mu - \psi \beta (\rho_1 P_D + \rho_2 (1 - P_D)). \quad (14)$$

The firm minimizes $F_A(\rho_1, \psi)$ when it gets an alarm from the IDS, and $F_N(\rho_2, \psi)$, when it does not get an alarm from the IDS. The user maximizes $H(\rho_1, \rho_2, \psi)$.

Table 3 Game in Strategic Form for the IDS Case

Firm's strategies	User's strategies	
	<i>H</i>	<i>NH</i>
(I, I)	$(c + (1 - \phi)d, \mu - \beta)$	$(c, 0)$
(I, NI)	$((c + (1 - \phi)d)P_D + d(1 - P_D), \mu - P_D\beta)$	$(cP_F, 0)$
(NI, I)	$(dP_D + (c + (1 - \phi)d)(1 - P_D), \mu - (1 - P_D)\beta)$	$(c(1 - P_F), 0)$
(NI, NI)	(d, μ)	$(0, 0)$

The following result holds in our model.

LEMMA 1. *Assuming that the IDS performs better than random guessing ($P_D > P_F$), the frequency of manual investigation is always higher in the alarm state than in the no-alarm state (i.e., $\rho_1 \geq \rho_2$). Additionally, the firm may conduct manual investigation in the no-alarm state only when it completely investigates all alarm states.*

Lemma 1 shows that the firm will investigate a larger fraction of users that generate alarms from the IDS, compared with those that do not generate alarms. The result also shows that the firm does not use mixed strategies in both alarm and no-alarm states. That is, if the probability of manual investigation in the alarm state is strictly less than one, then it does not investigate any user in the no-alarm state. If the probability of manual investigation in the alarm state is equal to one, then the firm may investigate a fraction of users that do not generate an alarm. The reason for this result is that the firm's expected benefit from manual investigation is greater in the alarm state than in the no-alarm state. Because the investigation is more efficient in the alarm state than in the no-alarm state, the firm expends all its investigation efforts in the alarm state before investing any investigation efforts in the no-alarm state. When the firm can no longer invest any more effort in the alarm state, i.e., when it investigates all cases that generate alarms, it starts to investigate in the no-alarm state. The equilibria for the IDS case are stated as the following result.

PROPOSITION 2. *The following mixed strategy profiles constitute the Nash equilibria for the IDS case.*

$$\begin{aligned} \text{If } \frac{\mu}{\beta} > P_D, \quad \text{then } & \left(\left(\rho_1 = 1, \rho_2 = \frac{\mu - P_D\beta}{(1 - P_D)\beta} \right), \right. \\ & \left. \psi = \frac{c(1 - P_F)}{c(P_D - P_F) + (1 - P_D)d\phi} \right). \\ \text{If } \frac{\mu}{\beta} \leq P_D, \quad \text{then } & \left(\left(\rho_1 = \frac{\mu}{P_D\beta}, \rho_2 = 0 \right), \right. \\ & \left. \psi = \frac{cP_F}{P_D d\phi - c(P_D - P_F)} \right). \end{aligned}$$

The IDS divides the parameter space into two distinct regions where different strategies are played, whereas both players play the same strategy throughout the parameter space in the no-IDS case. To further understand the effect of an IDS on the firm's

and users' strategies, we compare the no-IDS and IDS cases on key quantities of interest, such as hacking probability, investigation rate, and detection rate. The hacking probability, ψ , in the no-IDS and IDS cases, are already stated in Propositions 1 and 2, respectively. We define the investigation rate as the probability of investigating a user. The investigation rate for the no-IDS case is $\rho = \mu/\beta$ as given in Proposition 1. The investigation rate for the IDS case is given by $\rho_1 P(\text{alarm}) + \rho_2 P(\text{no-alarm})$, which is equal to

$$\left(\frac{\mu}{\beta} \right) \left(\frac{d\phi P_F}{d\phi P_D - c(P_D - P_F)} \right) \quad \text{if } \frac{\mu}{\beta} \leq P_D$$

and

$$\left(\frac{d\phi(\mu/\beta)(1 - P_F) - (d\phi - c)(P_D - P_F)}{d\phi(1 - P_D) + c(P_D - P_F)} \right) \quad \text{if } \frac{\mu}{\beta} > P_D.$$

The probability of detecting a hacker for the no-IDS case is given by $\rho (= \mu/\beta)$ in Proposition 1. For the IDS case, the detection rate is given by $\rho_1 P_D + \rho_2(1 - P_D)$, which is also computed to be equal to μ/β whether $\mu/\beta \leq P_D$ or $\mu/\beta > P_D$.

The following result compares the no-IDS and IDS cases.

PROPOSITION 3. (i) $\rho_2 \leq \rho \leq \rho_1$. *Compared to the no-IDS case,*

(ii) *The investigation rate is lower in the IDS case.*

(iii) *The probability of detecting a hacker is the same in the IDS case.*

(iv) *The hacking probability is higher (lower) in the IDS case if $\mu/\beta > P_D$ ($\mu/\beta \leq P_D$).*

The above result shows that, compared with the investigation rate in the no-IDS case, the investigation rate is higher for users that generate alarms from the IDS and lower for users that do not generate alarms from the IDS. However, the overall or effective investigation rate is lower in the IDS case compared with the no-IDS case, implying that an IDS enables the firm to better target its investigation of users. The effective detection rate is identical in the IDS and no-IDS cases. This result occurs because the firm's optimal strategy is to make the user indifferent between hacking and not hacking. The firm's strategy affects the user only through the detection rate, which determines the probability that a hacker will be detected. Consequently, the firm adjusts its strategy in the IDS

and no-IDS cases to keep the same level of detection rate.¹³ Though the detection rate is the same, the firm's strategies to maintain the same level of detection rate are different in the IDS and no-IDS cases. In case of IDS, the investigative resources are allocated first to the more efficient investigation (alarm state) and then to the less efficient one (no-alarm state). An interesting result is that although the probability of detecting a hacker is the same whether the firm does or does not use an IDS, the hacking probability is higher (lower) in the IDS case than in the no-IDS case if $\mu/\beta > P_D$ ($\mu/\beta \leq P_D$). To explain this, we should consider the user's optimal strategy in the equilibrium, which makes the firm indifferent between investigation and no investigation. In the IDS case, the user's hacking strategy and the IDS's probability of detection (P_D) affect the firm's cost. When P_D is high enough, i.e., when $\mu/\beta \leq P_D$, targeted investigation of users enabled by the IDS is so efficient in reducing the firm's cost that the user should reduce the hacking probability from that of the no-IDS case to make the firm indifferent between investigation and no-investigation. When P_D is not high enough, i.e., when $\mu/\beta > P_D$, the firm also has to extend its investigation to the no-alarm state. However, investigation in the no-alarm state is so inefficient¹⁴ in reducing the firm's cost that the user should increase the hacking probability from that of the no-IDS case to make the firm indifferent between investigation and no investigation.

Using Equation (13) and the results of Proposition 2, the firm's expected cost in the IDS case is

$$\frac{c}{\phi} \frac{1 - (\phi(1 - c/\phi d)P_D + (1 - \phi(1 - c/\phi d))P_F)}{1 - ((c/\phi d)P_F + (1 - c/\phi d)P_D)}$$

when $\mu/\beta > P_D$ and

$$\frac{c}{\phi} \frac{1}{c/\phi d + (1 - c/\phi d)P_D/P_F} \quad \text{when } \frac{\mu}{\beta} \leq P_D.$$

¹³ It follows from the mixed strategy equilibrium in the IDS and no-IDS cases. Formally, the firm sets $\mu - \rho\beta = \mu - (\text{DetectionRate}_{\text{No-IDS}})\beta = 0$ and $\mu - (\rho_1 P_D + \rho_2(1 - P_D))\beta = \mu - (\text{DetectionRate}_{\text{IDS}})\beta = 0$ in the no-IDS and IDS cases, respectively. Thus, $\text{DetectionRate}_{\text{No-IDS}} = \text{DetectionRate}_{\text{IDS}}$.

¹⁴ Investigations are less cost effective in the no-alarm state than in the no-IDS case because the likelihood that the user is a hacker is lower in the no-alarm state than in the no-IDS case.

5. The Value of IDS

In this section, we analyze if the firm should implement an IDS by deriving the value of an IDS. The firm will choose to implement an IDS if it offers a positive value to the firm. We calculate the value of IDS as (expected cost without the IDS) minus (expected cost with the IDS). In §4, we derived the firm's expected costs when it does and does not implement an IDS. For the IDS case, the cost was derived for a given IDS configuration that fixes the values of P_D and P_F . As stated in the introduction, firms often implement an IDS without configuring the IDS for their own environments, and some IDSs are nonconfigurable. In this section, we analyze whether the firm can benefit from the IDS when the IDS has an exogenously fixed configuration, and when the firm configures the IDS.

5.1. The Value of IDS with Default Configuration

The value of IDS is shown in Table 4. The positive value is what we would expect from an IDS. The most interesting finding occurs in the region where $\mu/\beta > P_D$; in this region, the use of an IDS increases the firm's cost. This result appears to reinforce concerns expressed by some security experts about IDS's value to firms. The finding that some firms may be hurt by the use of an IDS is particularly interesting despite the fact that the IDS provides useful information (better than random guesses) to the firm. The question that needs to be answered is as follows: Why is the value positive when $\mu/\beta \leq P_D$ and negative when $\mu/\beta > P_D$? The answer to this question lies in the hacking probability derived in Proposition 3. We explain the value of IDS as follows: The firm's loss is composed of two components—investigation cost and expected damage. The investigation cost is increasing in the investigation rate, and the damage cost is increasing in the hacking probability. As noted in Proposition 3, when $\mu/\beta \leq P_D$, the investigation rate and hacking probability are lower in the IDS case

Table 4 The Value of IDS

Regions	The value of IDS	Is an IDS beneficial?
$\frac{\mu}{\beta} > P_D$	$-\frac{c}{\phi} \frac{(P_D - P_F)(1 - \phi)(d\phi - c)}{d\phi - ((d\phi - c)P_D + cP_F)}$	No
$\frac{\mu}{\beta} \leq P_D$	$\frac{c}{\phi} \frac{(P_D - P_F)(d\phi - c)}{P_D d\phi - c(P_D - P_F)}$	Yes

than in the no-IDS case. The detection rate remains the same. Consequently, both investigation cost and damage cost are lower in the IDS case than in the no-IDS case. When $\mu/\beta > P_D$, the investigation rate is lower, and consequently the investigation cost is also lower in the IDS case than in the no-IDS case. However, a higher hacking probability in the IDS case compared with the no-IDS case (as shown in Proposition 3) results in a higher damage cost in the IDS case. The higher damage cost more than offsets the savings realized in investigation cost, causing the total cost to be higher in the IDS case than in the no-IDS case.

An IDS deters a hacker if the probability of hacking is lower when the IDS is deployed than when it is not. We find that the IDS is valuable only when the IDS deters users from hacking. This finding supports the common belief that the goal of an IDS is not only to detect intrusions, but also to act as a deterrent against intrusions (Fisch and White 2000, p. 90). The following statements from the National Institute of Standards and Technology's (NIST's) special publication on IDSs (Bace and Mell 2001) also articulate this belief: "A fundamental goal of security management is to affect the behavior of individual users in a way that protects information systems from security problems. Intrusion detection systems help organizations accomplish this goal by increasing the perceived risk of discovery and punishment of attackers. This serves as a significant deterrent to those who violate security policy" (p. 6). Our finding strongly supports the above rationale for using an IDS, and makes an even stronger case that an IDS will be of value only when the IDS deters hackers. Propositions 3 and 4 show that the value of IDSs should be assessed, not simply as the benefit derived from improved detection of intrusions, but also from increased deterrence of hackers.

The following proposition summarizes our finding about the value of IDS when the configuration is assumed to be exogenous.

PROPOSITION 4. *For the default configuration case, (i) The value of IDS is nonnegative when $\mu/\beta \leq P_D$ and negative when $\mu/\beta > P_D$; (ii) the value of IDS is positive only if the IDS is a deterrent to hackers.*

Proposition 4 cautions firms against using the default configuration, because it may be detrimental

to firms. Note that when the detection rate of the IDS, P_D , is smaller than μ/β , the value of IDS is negative. If the IDS is nonconfigurable, then the firm will not use an IDS if $\mu/\beta > P_D$. However, if the IDS is configurable, then the firm faces the question of what value of P_D (and P_F) it should choose and whether the firm will always be better off with an IDS. We analyze this case next.

5.2. The Value of IDS with Optimal Configuration

Proposition 2 shows that the firm can be in one of two regions with the default configuration. By choosing P_D (or configuring the IDS), the firm can determine the region where the firm will operate. A comparison of costs in the two equilibrium regions in the IDS case shows that the firm realizes a lower cost when $\mu/\beta \leq P_D$. Consequently, the firm will choose the value of P_D so that this condition is satisfied. Next the firm should decide where to lie within this region. Writing the cost when $\mu/\beta \leq P_D$ as a function of P_D and taking the first derivative gives

$$\frac{\partial(\cdot)}{\partial P_D} = \frac{cdP_D^{1/r}}{(d\phi - c)P_D + cP_D^{1/r}} \geq 0. \quad (15)$$

This derivative implies that the firm will choose to set P_D as small as possible. Because it is optimal for the firm to be in Region 1, the firm sets P_D of its IDS to μ/β . That is, the optimal configuration for the IDS is

$$P_D = \frac{\mu}{\beta} \quad (16)$$

$$P_F = \left[\frac{\mu}{\beta} \right]^{1/r}. \quad (17)$$

Substituting the above optimal configuration point into the cost expression gives an expected cost for the firm of

$$\frac{c}{\phi} \frac{d\phi}{(\mu/\beta)^{1-1/r} d\phi + (1 - (\mu/\beta)^{1-1/r})c}.$$

The value of an optimally configured IDS can be calculated to be

$$\frac{c}{\phi} \left(1 - \frac{d\phi}{(\mu/\beta)^{1-1/r} d\phi + (1 - (\mu/\beta)^{1-1/r})c} \right),$$

which is nonnegative.

The following result holds for an optimally configured IDS.

PROPOSITION 5. (i) *The value of the optimally configured IDS is nonnegative.*

(ii) *The optimally configured IDS deters hackers.*

(iii) *The optimally configured IDS yields the same investigation strategy as a perfect ($P_D = 1$ and $P_F = 0$) IDS.*

Proposition 5 is in contrast to the result for the default configuration, which showed that the use of an IDS will be detrimental to firms if $P_D < \mu/\beta$. Proposition 5 provides strong theoretical support to guidelines that warn firms against using out-of-box configuration for IDSs (McHugh 2000, Amoroso 1999, Mell et al. 2002). Furthermore, the firm configures the IDS such that it will investigate all users that generated alarms from the IDS ($\rho_1 = 1$) and none of the users that did not generate alarms ($\rho_2 = 0$). This investigation strategy is also optimal for the case when the IDS is perfect ($P_D = 1$ and $P_F = 0$). That is, an optimally configured imperfect IDS yields the same manual investigation strategy as the perfect IDS.

5.3. Analysis of the Value of IDS

To derive further insights into the effects of firm, user, and technology parameters on the value of IDS as well as other quantities of interest, such as the hacking rate, investigation rate, and detection rate, we performed comparative static analysis of these quantities for the optimally configured IDS. The exact expressions for this analysis are available from the authors. Table 5 shows the directions of various effects. We discuss these results in the following paragraphs.

Table 5 Comparative Statics for the Optimally Configured IDS

Panel A The Effect of Model Parameters on Hacking, Investigation, and Detection Rates						
	$c/d\phi$	μ/β			r	
Hacking rate	+	+			+	
Investigation rate	+	+			+	
Detection rate	0	+			0	

Panel B The Effect of Model Parameters on the Value of IDS						
	c	d	β	μ	ϕ	r
Value of IDS	- if $c < a$ + if $c > a$	+	+	-	- if $\phi > b$ + if $\phi < b$	-

Note. $a = \frac{d\phi(\mu/\beta)^{(1/2)(1-1/r)}}{(\mu/\beta)^{(1/2)(1-1/r)} - 1}$, $b = \frac{c}{d} \left(\frac{1 - (\mu/\beta)^{(1-r)/r}}{1 - (\mu/\beta)^{(1-r)/(2r)}} \right)$.


Effects on Hacking and Investigation Rates. The hacking rate and the investigation rate are increasing in $c/(d\phi)$ and μ/β . A higher value for $c/(d\phi)$ implies a less efficient manual investigation for the firm, and a higher value for μ/β implies a higher expected benefit to the user from hacking. An increase in either of these quantities offers more incentives to users to hack, which in turn causes the firm to investigate more frequently. An increase in r also increases the hacking and investigation rates. Because a higher r implies a lower-quality IDS technology (i.e., less precise targeting of users), the users increase hacking, and the firm increases its investigation rate to compensate for the lower quality.

Effect on Detection Rate. The effective detection rate is equal to μ/β in the IDS and no-IDS cases. Consequently, the firm's parameters do not have any effect on the detection rate. A higher (lower) μ/β results in a higher (lower) detection rate.

Effects on IDS Value. The value of IDS is decreasing in μ/β and r . A lower quality of IDS technology reducing its value is not surprising. However, an increase in μ/β , which increases the incentives of users to hack, reducing the value obtained from an IDS, is counterintuitive. This result can be explained as follows. A higher μ/β causes the firm to configure the IDS to work at a higher P_D (note that $P_D = \mu/\beta$ at the optimal configuration) and a higher P_F . Because the firm investigates all users that generate alarms, a higher rate of false positives increases the firm's investigation cost. Consequently, when μ/β increases, the higher level of hacking, combined with a higher level of investigation, reduces the value from the IDS.

Unlike the hacking and investigation rates, the effects of firm parameters such as c , d , and ϕ on the IDS value depend on their individual values rather than on $c/(d\phi)$. A higher d increases the IDS value, a result that is to be expected. An increase in c or ϕ may either increase or decrease the IDS value. Figures 4a and 4b show the regions where the effects are positive and negative. An increase in c increases the IDS value in the region characterized by low $c/(d\phi)$ and μ/β values. This region offers a low incentive for users to hack. That is, in the low hacking region an IDS is more valuable when the manual investigation cost is higher. An increase in c reduces the IDS value in the region characterized by high $c/(d\phi)$ and μ/β

Figure 4 The Effect of (a) Inspection Cost and (b) Recovery Rate on the Value of IDS



values. This region offers a high incentive for users to hack. The effect of ϕ is qualitatively the opposite of the effect of c .

6. Managerial Implications, Discussions, and Extensions

6.1. Managerial Implications and Discussions

Deterrence is the key to the value of IDS. The benefit of deploying an IDS depends on how much it prevents hackers from committing intrusions. Although IDSs are classified as detective controls because they detect attacks that were not prevented, they implicitly act as preventive controls by changing the behavior of attackers in the first place, and thus eliminating attacks. According to Proctor (2001), “The presence of a network-based IDS can put [external] hackers on notice that their actions may lead to legal action” (p. 43). “Host-based systems provide very similar deterrent effect [to internal hackers]. People who know that their actions may be monitored are less likely to commit misuse. I consider this to be the greatest benefit. . . . Unfortunately it is very difficult to quantify” (p. 70). Our results reinforce those claims. We not only quantify the benefits, but also demonstrate that the deterrence aspect of IDSs makes them valuable.

Out-of-box implementations have been highlighted as the biggest impediment to deriving value from security controls in general (McCarthy 1998). Within an intrusion detection context, firms have been warned not to implement an IDS as a plug-and-play black-box solution (Amoroso 1999, p. 17). Our analysis revealed that optimally configured IDSs always provide nonnegative value to their adopters. Optimal configuration requires user-specific parameters

and the exact shape of the ROC curve. Obtaining these data may require additional effort from firms before implementing the IDS. By using the out-of-box configuration, firms may be taking the easy way out, but they may be hurting themselves. This might explain why some firms have discontinued the use of IDSs after realizing that IDSs did not offer any value (Panko 2003, p. 381) and the conclusions of Gartner Inc. (Gartner 2003) about the failure of IDS technology. We find that IDSs are not inherently useless, as depicted in the Gartner report, but that they lose their true potential if not configured properly.

Another current widespread complaint against IDSs is that they produce many false alarms: “False positives are tremendous time wasters and drive up operational labor costs. They also create so much noise that they drown out security-critical events that truly require attention” (Proctor 2001, p. 108). Although this complaint is perfectly legitimate, we show that this limitation of current IDSs does not make them useless. Some firms, especially those in high hacking environments, should choose to operate at a high false alarm rate, higher than the default false alarm rate, to get a detection rate high enough to benefit from their IDS implementations. Our research points out that reducing the false alarm error rate is not always the best strategy.

A too-tight configuration, characterized by a low detection rate, $P_D < \mu/\beta$, yields a low number of false alarms, but increases the firm’s cost. The firm is better off not using the IDS in this case. If the configuration is too loose, $P_D > \mu/\beta$, while the firm does realize positive value from the IDS, the high false alarm rate prevents the firm from realizing the full value of the IDS. The optimal configuration, $P_D = \mu/\beta$, depends on the user parameters or the external environment in which the firm is operating. Although the firm’s cost parameters affect the value of IDS, they do not affect the optimal configuration. The firm’s parameters also do not affect whether the firm realizes a positive value from an IDS. These findings point to the need for assessing the external environment properly before using an IDS. Current guidelines on implementation of IDSs in security architecture emphasize the firm’s operational parameters, such as costs (CERT 2000). Our results suggest that the hackers’ incentives (external environment) play a more significant role

than firms' costs in determining whether or not an IDS benefits firms.

IDS developers should also pay close attention to the configuration issue. They should design IDSs that are easy to configure, especially in light of high false positive rates associated with IDSs. Panko (2003, p. 381) observes that "Tuning is so essential that one might assume that IDSs make it easy to do: however, most IDSs today make tuning extremely difficult. A high number of false positives, coupled with the difficulty of tuning, have caused many firms to avoid IDSs or to remove already-implemented IDSs." Therefore, IDS vendors should design products that can be easily configured to fit the operating environment.

Our findings suggest that the IDS developers should also focus on the environment in which the deploying firm operates. The default setting for the IDS should have a high (low) P_D value for an environment where the users have high (low) incentives to hack. Some industries, such as defense and financial institutions, may offer more incentives for users to hack, and others may offer fewer incentives. The IDS developers should calibrate their products differently for different industries.

The most important performance-related data of an IDS is its ROC curve. Most vendors do not provide these data. Several groups, including academic institutions, research laboratories, and commercial organizations, have tested commercial and government-sponsored IDS products (Debar et al. 1998; Aguirre and Hill 1997; Puketza et al. 1997; Lippmann et al. 2000a, b; Durst et al. 1999; Mueller and Shipley 2001; NSS Group 2001; Yocom and Brown 2001). Only a few of these tests included false positive and false negative rates as performance metrics (Lippmann et al. 2000a, b). Although these evaluations are valuable attempts to understand the quality of IDSs (McHugh 2000), the value of results from these evaluations may be limited, because these tests primarily used default configurations with no tuning relative to the environment in which devices were employed. Rigorous testing of IDSs at various operating points on the ROC curve is necessary for firms to realize the full potential of IDSs.

6.2. Extensions

In this section, we relax several of the assumptions made in our basic model and derive the effects of these on our results.

Risk Neutrality Assumption. We performed our analysis by assuming that the firm and users are risk neutral. Consequently, utility was assumed to be a linear function of benefits. Both the firm and users could have other risk dispositions. For example, a firm could be risk averse with respect to critical IT assets, or a user could be a risk seeker. We also investigated the effects of players' risk modes on our results. For this analysis, we let $U_u(x)$ and $U_f(x)$ —where x represents payoff—be the utility functions of users and the firm, respectively. U is increasing in x , and $U(x)$ has the same sign as x . For a risk-averse (-seeking) player, U is concave (convex). We did not change any other aspect of our basic model.

In the no-IDS case, the firm and a user maximize their expected utilities given by

$$F(\rho, \psi) = \rho U_f(-c) + \rho \psi (U_f(-c - (1 - \phi)d) - U_f(-c)) \\ + \psi (1 - \rho) U_f(-d) \\ H(\rho, \psi) = \psi U_u(\mu) + \rho \psi (U_u(\mu - \beta) - U_u(\mu)),$$

respectively. The equilibrium is given by the mixed strategy profile

$$\left(\rho = \frac{U_u(\mu)}{U_u(\mu) - U_u(\mu - \beta)}, \right. \\ \left. \psi = \frac{U_f(-c)}{U_f(-c) + U_f(-d) - U_f(-c - (1 - \phi)d)} \right).$$

The firm's expected utility at the equilibrium is equal to

$$\frac{U_f(-c)U_f(-d)}{U_f(-c) + U_f(-d) - U_f(-c - (1 - \phi)d)}.$$

In the IDS case, the firm's expected utility in the alarm and no-alarm states are given by

$$F_A(\rho_1, \rho_2, \psi) = \rho_1 U_f(-c) + \rho_1 \eta_1 (U_f(-c - (1 - \phi)d) \\ - U_f(-c)) + \eta_1 (1 - \rho_1) U_f(-d) \\ F_N(\rho_1, \rho_2, \psi) = \rho_2 U_f(-c) + \rho_2 \eta_2 (U_f(-c - (1 - \phi)d) \\ - U_f(-c)) + \eta_2 (1 - \rho_2) U_f(-d),$$

respectively. A user's expected utility is

$$H(\rho_1, \rho_2, \psi) = \psi U_u(\mu) + \psi(\rho_1 P_D + \rho_2(1 - P_D)) \cdot (U_u(\mu - \beta) - U_u(\mu)).$$

The equilibrium is given by

If $\frac{U_u(\mu)}{U_u(\mu) - U_u(\mu - \beta)} > P_D$, then

$$\left(\left(\rho_1 = 1, \rho_2 = \frac{U_u(\mu) - P_D(U_u(\mu) - U_u(\mu - \beta))}{(1 - P_D)(U_u(\mu) - U_u(\mu - \beta))} \right), \right. \\ \left. \psi = \frac{U_f(-c)(1 - P_F)}{(1 - P_D)(U_f(-d) - U_f(-c - (1 - \phi)d)) + U_f(-c)(1 - P_F)} \right).$$

If $\frac{U_u(\mu)}{U_u(\mu) - U_u(\mu - \beta)} \leq P_D$, then

$$\left(\left(\rho_1 = \frac{U_u(\mu)}{P_D(U_u(\mu) - U_u(\mu - \beta))}, \rho_2 = 0 \right), \right. \\ \left. \psi = \frac{U_f(-c)P_F}{P_D(U_f(-d) - U_f(-c - (1 - \phi)d)) + U_f(-c)P_F} \right).$$

We show that Proposition 3 and Proposition 4 hold for any $U_u(x)$ and $U_f(x)$. Our analysis shows that the risk dispositions of players do not affect our results qualitatively. However, for a given set of parameter values, the risk profiles do affect the region where the game will be played, and consequently whether the firm will realize a positive or negative value from the deployment of an IDS.

The Exponential Distribution Assumption. We derived the ROC curve function by assuming that the numerical score computed by the IDS to distinguish hackers and normal users followed an exponential distribution. The exponential distribution assumption yielded the ROC curve given by $P_D = P_F'$. We used this functional relationship in our analysis to derive the optimal configuration of IDS. Specifically, it was used to derive Equation (15). The expressions for the value of IDS and the analysis of IDS value for the out-of-box IDS configuration remain the same for any functional relationship between P_D and P_F .

If we do not make any distributional assumption, then P_D and P_F are implicitly related, through t , by the following equations:

$$P_D = \int_t^\infty f_H(x) dx, \\ P_F = \int_t^\infty f_N(x) dx.$$

An analysis of the value of IDS when $\mu/\beta \leq P_D$ shows that $P_D = \mu/\beta$ remains optimal if

$$\frac{f_H(t)}{1 - F_H(t)} < \frac{f_N(t)}{1 - F_N(t)} \cdot \frac{f(x)}{1 - F(x)}$$

is the hazard rate of the distribution. The hazard rate is increasing in x for many distributions including the uniform, the normal, the Pareto, the logistic, the exponential, and any distribution with nondecreasing density. For a distribution that has an increasing hazard rate, if the probability distributions for the normal users and hackers differ only in their mean values, then we can show that

$$\frac{f_H(t)}{1 - F_H(t)} < \frac{f_N(t)}{1 - F_N(t)},$$

and consequently our results hold for all the above-mentioned distributions.

Relationship Among c , d , and ϕ . We assumed that the cost of manual investigation, the damage from an undetected intrusion, and the fraction of the damage prevented or recovered when a manual investigation detects an intrusion are independent of each other. However, functional relationships may exist among these parameters. For example, the cost of investigation and the fraction of damage recovered may be positively related, i.e., an increase in c may increase ϕ , and vice versa. A similar relationship may also exist between c and d as well as ϕ and d . We analyzed the effect of such dependencies on our results by assuming that c is a function of ϕ and d , i.e., $c(\phi, d)$.¹⁵ The analysis shows that the equilibria for the no-IDS and IDS cases stated as Proposition 1 and Proposition 2, respectively, change only to the extent of changing the parameter c to function $c(\phi, d)$ in these equilibria. Lemma 1 and Proposition 3 hold, even under the new assumption of $c(\phi, d)$. Whereas the expressions for the value of IDS given in Table 4 change to reflect the assumption, the result that the value is negative (nonnegative) when $\mu/\beta > P_D$ ($\mu/\beta \leq P_D$) does not change. Consequently, the firm will still set

¹⁵ We could have also modeled the dependence between ϕ and d as $\phi(d)$. However, the results do not change qualitatively because, as discussed later in this section, the results depend on the value of $c/(d\phi)$. Dependency among these parameters affects our results only to the extent that $c/(d\phi)$ will be replaced by the expression that captures the functional relationship between ϕ and d .

P_D such that $\mu/\beta \leq P_D$. That is, Proposition 4 holds under $c(\phi, d)$. The optimal configuration parameters also do not change. The primary reason that all our results remain valid for any form of functional dependency among c , d , and ϕ is that most of our analytical results can be stated as functions of $c/(d\phi)$. That is, any dependency among these parameters replaces $c/(d\phi)$ by the expression that models the dependency.

We had also assumed that c , d , and ϕ are exogenous parameters, and the firm's decisions are to (i) decide whether it should deploy an IDS, (ii) choose the IDS configuration, if an IDS is deployed, and (iii) decide its manual investigation strategy. That is, we assumed that the operating environment of the firm, defined by its cost parameters, c , d , and ϕ , is exogenously given and focused on decisions related to the implementation of the IDS technology. A model that allows the firm to choose its operating environment, in addition to the technology aspects, will offer us insights into the interaction between decisions concerning the selection of technology and operating environments.

Similar to the dependency among c , d , and ϕ , several other dependencies among exogenous parameters, such as dependency among μ , β , d , and ϕ , do not change the results of the paper qualitatively.

7. Conclusions, Limitations, and Future Research Directions

7.1. Conclusions

IT security management addresses three fundamental components of security: prevention, detection, and response. All these factors are indispensable parts of effective security programs, and, therefore, should be carefully designed and deployed. However, firms have traditionally emphasized prevention over detection and response. After all, if threats are prevented detection and response are unnecessary. Recently, organizations have realized that it is impossible to eliminate all security risks. As a result, detection-based systems have started to gain popularity in the IT security domain. Today, IDSs are the most popular detective controls. Although IDS has been the fastest-growing security product in the market for the last few years, the security community is uncertain

about their value. Our research was aimed at providing insights into the value of these mechanisms.

We showed that a firm might not realize a positive value from an improperly configured IDS. An improperly configured IDS may encourage more hacking, resulting in a higher loss for the firm. An optimally configured IDS deters hackers, thus, reducing the need for investigation by security experts for security violations. We also showed that optimal configuration depends not on the firm's internal cost parameters, but on the external hacker parameters. This result highlights the significance of understanding hacker behavior and motivation when employing an IDS. The findings of our research shed light on concerns regarding the value of IDSs. Our results should be reassuring to firms that have implemented a properly configured IDS. To firms that are using default configuration or that have not adopted an IDS because of doubts about its value, our results provide incentives to implement an appropriately configured IDS.

7.2. Limitations and Future Research Directions

We made certain simplifying assumptions in our basic model to analyze the value of IDSs. Relaxation of many of these assumptions does not change the qualitative nature of our results, as shown in §6.2. As with all models, our model has a number of limitations. In this paper, we assumed that the model parameters were common knowledge to the firm and users. One area that seems particularly interesting is games with incomplete information, in which either the firm or the user is uncertain about the other's payoffs. This perspective allows incorporation of uncertainty about the nature of the game being played.¹⁶ If the firm is uncertain about the utility of hacking to the user, it may use its prior knowledge to develop a monitoring strategy. For instance, assume that the firm believes that a user's utility can be high or low. We can represent the firm's beliefs with a subjective probability distribution, p_H and p_L . In other words, the firm places a $100p_H\%$ chance that it is dealing with a high utility user and a $100p_L\%$ chance that it is dealing with a low

¹⁶ Many of these common knowledge related assumptions have been analyzed by game theorists. Hansanyi (1967, 1968a, b) has shown that the lack of knowledge about payoffs does not alter the basic representation of a game and the qualitative nature of the results.

utility user. However, the user knows his true utility. We can analyze this setting using Harsanyi transformation, in which nature makes the first move and chooses the user type in accordance with the firm's subjective probabilities. We leave the detailed analysis of this model to future research.

We also assumed a one-shot setting in which the firm makes its configuration and investigation probabilities, and the user makes its hacking decision. However, it may be more realistic to consider a multiperiod model in which the firm revises its estimates every period based on its observations of the hacker's strategy in previous periods. Such learning has been analyzed in game theory. Fudenberg and Levine (1998, pp. 34–35, Prop. 2.2) show that when players learn but use a myopic approach every period, if the empirical distribution over each player's choices converges, then the strategy profile is a Nash equilibrium. If convergence is achieved, then the dynamic model and the static game theory-based model yield identical outcomes in the equilibrium. Fudenberg and Levine also state that the empirical distributions need not always converge. In addition, the type of learning model employed also has an effect on the convergence. For instance, an open issue is what type of learning model is appropriate for our context. Some of the questions include the following. Is learning based only on the most recent move, or is it based on the history of all moves? What relative weights should be assigned to different moves? How are the probabilities updated based on the history? Answers to these questions will provide valuable additional insights into the trade-off between a static model and a dynamic model. A valuable extension of our research is to analyze a dynamic model that incorporates learning, and to compare the results of this model with the static game theory approach presented in this paper.

In our analysis, we assumed that security experts take appropriate actions after receiving alarms from IDSs. This approach, also called passive response, is the current trend in commercial IDSs. Another response option is to let the IDS take an action without human intervention (active response).¹⁷ Current

IDSs provide little or no guidance to security management once an attack has been identified (Allen et al. 2000). However, if IDSs provide information about the type of attack and how to handle it, the firm may be able to increase the fraction of damage recovered or reduce the manual investigation effort, which in turn can lead to a higher value from IDSs. An extension of our work could address the effect of guidance of the IDS on recovery and response.

Notwithstanding these potentially attractive avenues for further research, the present study, which we believe is one of the first that investigates the value of specific IT security technologies, provides useful insights into value from and configuration of IDSs.

References

- Aguirre, S. J., W. H. Hill. 1997. Intrusion detection fly-off: Implications for the United States Navy. MITRE Technical Report MTR 97W096, McLean, VA.
- Allen, J., A. Christie, W. Fithen, J. McHugh, J. Pickel, E. Stoner. 2000. State of the practice of intrusion detection technologies. Technical Report CMU/SEI-99-TR-028 ESC-99-028, Pittsburgh, PA.
- Alpert, B. 1999. As e-tailing booms on the net, so does the demand for virtual security. *Barron's* 79(4) 25.
- Amoroso, E. 1999. *Intrusion Detection*. Intrusion.Net Books, NJ.
- Axelsson, S. 2000. The base-rate fallacy and the difficulty of intrusion detection. *ACM Trans. Inform. System Security* 3(3) 186–205.
- Bace, R., P. Mell. 2001. NIST special publication on intrusion detection systems. SP-800-31. National Institute of Standards and Technology.
- Baiman, S. 1982. Agency research in managerial accounting: A survey. *J. Accounting Literature* 1 154–213.
- Bashir, I, E. Serafini, K. Wall. 2001. Securing network software applications: Introduction. *Comm. ACM* 44(2) 28–30.
- Becker, G. 1968. Crime and punishment: An economic approach. *J. Political Econom.* 76 169–217.
- CERT (Computer Emergency and Response Team). 2000. *Detecting Signs of Intrusion*. CERT Security Improvement Modules, Pittsburgh, PA.
- Debar, H., M. Dacier, A. Wespi, S. Lampart. 1998. *A Workbench for Intrusion Detection Systems*. IBM Zurich Laboratory, Ruschlikon, Switzerland.
- Denning, D. E. 1987. An intrusion detection model. *IEEE Trans. Software Engrg.* 13(2) 222–232.
- Denning, D. E. 2000. Reflections on cyberweapons controls. *Comput. Security J.* 16(4) 43–53.

¹⁷ Possible active responses are terminating the network session by resetting the TCP connection or updating the firewall rule set to

block packets coming from the IP address that appears to be the source of the intrusion. However, attackers might use IP-spoofing tools to trick the firewall in order to get into the system. Blocking the traffic may also result in denial of service (DOS) (Proctor 2001).

- D'haeseleer, P., S. Forrest, P. Helman. 1996. An immunological approach to change detection: Algorithms, analysis, and implications. *Proc. IEEE Sympos. Security Privacy*, 110–119.
- Durst, R., T. Champion, B. Witten, E. Miller, L. Spagnuolo. 1999. Testing and evaluating computer intrusion detection systems. *Comm. ACM* 42(7) 53–61.
- Dye, R. A. 1986. Optimal monitoring policies in agencies. *RAND J. Econom.* 17 339–350.
- Escamilla, T. 1998. *Intrusion Detection: Network Security Beyond the Firewall*. John Wiley & Sons, New York.
- Feichtinger, G. 1983. A differential games solution to a model of competition between a thief and the police. *Management Sci.* 29 686–699.
- Fellingham, J. C., P. Newman. 1985. Strategic considerations in auditing. *Accounting Rev.* 60(4) 634–650.
- Fisch, E. A., G. B. White. 2000. *Secure Computer and Networks*. CRC Press, Boca Raton, FL.
- Frincke, D., J. Evans, D. Aucutt. 1996. Hierarchical management of misuse reports. *Proc. Internat. Conf. Comput. Inform.*, Ontario, Canada.
- Fudenberg, D., D. Levine. 1998. *The Theory of Learning in Games*. MIT Press, Cambridge, MA.
- Fudenberg, D., J. Tirole. 1993. *Game Theory*. MIT Press, Cambridge, MA.
- Gartner. 2003. Hype cycle for information security. Gartner Research Report (May 30), Stamford, CT.
- Garvey, T., T. Lunt. 1991. Model-based intrusion detection. *Proc. 14th National Comput. Security Conf.*, Washington, D.C.
- Goodman, M. D., S. W. Brenner. 2002. The emerging consensus on criminal conduct in cyberspace. *UCLA J. Law Tech.* (3).
- Hansanyi, J. C. 1967. Games with incomplete information played by Bayesian players, I: Basic model. *Management Sci.* 14(3) 159–182.
- Hansanyi, J. C. 1968a. Games with incomplete information played by Bayesian players, II: Bayesian equilibrium points. *Management Sci.* 14(5) 320–334.
- Hansanyi, J. C. 1968b. Games with incomplete information played by Bayesian players, III: The basic probability distribution of the game. *Management Sci.* 14(7) 486–502.
- Hulme, H. 2002. Businesses keep spending on security. *Inform. Week* (January 28).
- Ilgun, K. 1992. Ustat: A real-time intrusion detection system for Unix. Master's thesis, Computer Science Department, University of California at Santa Barbara, CA.
- Internet Security Systems. 2001. The truth about false positives. Technical White Paper, Internet Security Systems, Atlanta, GA.
- Jajodia, S., J. Miller. 1993. Editor's preface. *J. Comput. Security* 16(4) 43–53.
- Kanodia, C. S. 1985. Stochastic and moral hazard. *J. Accounting Res.* 23 175–193.
- Kilgour, D. M. 1992. Site selection for on-site inspection in arms control. *Arms Control* 13(13) 439–462.
- Koerner, B. I. 1999. Who are hackers, anyway? *U.S. News World Rep.* 17(2) 53.
- Kumar, S., E. Spafford. 1996. A pattern matching model for misuse intrusion detection. *The COAST Project*. Purdue University, West Lafayette, IN.
- Lee, W., W. Fan, M. Miller, S. Stolfo, E. Zadok. 2002. Toward cost-sensitive modeling for intrusion detection and response. *J. Comput. Security* 10(1/2) 5–22.
- Lippmann, R. P., D. J. Fried, I. Graf, J. W. Haines, K. R. Kendall, D. McClung, D. Weber, S. E. Webster, D. Wyschogrod, R. K. Cunningham, M. A. Zissman. 2000a. Evaluating intrusion detection systems: The 1998 DARPA off-line intrusion detection evaluation. *Proc. 2000 DARPA Inform. Survivability Conf. Exposition (DISCEX)* 2 12–26.
- Lippmann, R. P., J. W. Haines, D. J. Fried, J. Kobra, K. Das. 2000b. The 1999 DARPA off-line intrusion detection evaluation. *Comput. Networks* 34(2) 579–595.
- Lunt, T. 1990. Ides: An intelligent system for detecting intruders. *Proc. Sympos.: Comput. Security, Threat Countermeasures*, Rome, Italy.
- Lunt, T. 1993. A survey of intrusion detection systems. *Comput. Security* 12 405–418.
- Lunt, T., R. Jagannathan. 1988. A prototype real-time intrusion detection system. *Proc. 1988 IEEE Sympos. Security Privacy*, Oakland, CA.
- Lunt, T., A. Tamaru, F. Gilham, R. Jagannathan, R. C. Jalali, H. Javitz, A. Valdos, P. Neumann, T. Garvey. 1992. A real-time intrusion detection expert system. Technical report, Consumer Science Laboratory, SRI International, Menlo Park, CA.
- Maschler, M. 1966. A price leadership method for solving the inspector's non-constant-sum game. *Naval Res. Logist. Quart.* 13 11–33.
- Maschler, M. 1967. The inspector's non-constant-sum game: Its dependence on a system of detectors. *Naval Res. Logist. Quart.* 14 275–290.
- McCarthy, L. 1998. *Intranet Security*. Sun Microsystems Press, Santa Clara, CA.
- McHugh, J. 2000. Testing intrusion detection systems: A critique of the 1998 and 1999 DARPA intrusion detection system evaluations as performed by Lincoln Laboratory. *ACM Trans. Inform. System Security* 3(4) 262–294.
- McHugh, J., A. C. Christie, J. Allen. 2000. Defending yourself: The role of intrusion detection systems. *IEEE Software* 17(5) 42–51.
- Mell, P., V. Hu, R. Lippmann, J. Haines, M. Zissman. 2002. *An Overview of Issues in Testing Intrusion Detection Systems*. NIST IR 7007, Gaithersburg, MD.
- Messmer, E. 1999. Getting the drop on network intruders. *Network World* (October 4).
- Mishra, B. K., P. Newman, C. Stinson. 1997. Environmental regulations and incentives for compliance audits. *J. Accounting Public Policy* 16(2) 187–214.
- Monrose, F., A. Rubin. 1997. Authentication via keystroke dynamics. *4th ACM Conf. Comput. Comm. Security*, Zurich, Switzerland.
- Mookherjee, D., I. P. L. Png. 1992. Monitoring vis-à-vis investigation in enforcement of law. *Amer. Econom. Rev.* 82(3) 556–565.
- Mueller, P., G. Shipley. 2001. Dragon claws its way to the top. *Network Comput.* 20(August) 45–67.
- National Computer Security Center. 1988. *A Guide to Understanding Audit in Trusted Systems*. NCSC-TG-001, Version 2 (June), The Rainbow Series, Meade, MD.
- Neumann, P., P. Porras. 1999. Experience with emerald to date. *Proc. 1st USENIX Workshop Intrusion Detection Network Monitoring*, Santa Clara, CA, 73–80.
- Newman, P., S. Rhoades, R. Smith. 1996. Allocating audit resources to detect fraud. *Rev. Accounting Stud.* 1 161–182.
- NIST Publication 800-12. 1996. *An Introduction to Computer Security*. National Institute of Standards and Technology, Gaithersburg, MD.
- Northcutt, S. 1999. Evaluating intrusion detection systems without attacking your friends. *Network Intrusion Detection* 86.

- NSS Group. 2001. *Intrusion Detection Systems Group Test*, Ed. 2. Oakwood House, Wennington, Cambridgeshire, UK (December).
- Panko, R. 2003. *Corporate Computer and Network Security*. Prentice Hall, NJ.
- Peltier, T. R. 2001. *Information Security Risk Analysis*. Auerbach Publications, Boca Raton, FL.
- Polinsky, A., S. Shavell. 1979. The optimal trade-off between the probability and magnitude of fines. *Amer. Econom. Rev.* **69** 880–891.
- Porras, P., R. Kemmerer. 1992. Penetration state transition analysis: A rule-based intrusion detection approach. *IEEE 8th Annual Comput. Security Appl. Conf.*, San Antonio, TX, 220–229.
- Porras, P., P. Neumann. 1997. Emerald: Event monitoring enabling responses to anomalous live disturbances. *Proc. 20th Nat. Inform. Systems Security Conf.*, Baltimore, MD, 353–365.
- Power, R. 2002. CSI/FBI computer crime and security survey. *Comput. Security Issues Trends* **8**(1) 1–22.
- Proctor, P. E. 2001. *The Practical Intrusion Detection Handbook*. Prentice Hall, NJ.
- Ptacek, T. H., T. N. Newsham. 1998. *Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection*. Secure Networks Inc., Calgary, Alberta, Canada.
- Puketza, N., M. Chung, R. O. Olsson, B. Mukherjee. 1997. A software platform for testing intrusion detection systems. *IEEE Software* **14**(5) 43–51.
- Rothke, B. 2000. Hackers then and now: Answers to some perennial questions. *Comput. Security J.* **16**(3) 11–14.
- Russell, D., G. T. Gangemi. 1992. *Computer Security Basics*. O'Reilly & Associates, Inc., Sebastopol, CA.
- Russell, G. S. 1990. Game models for structuring monitoring and enforcement systems. *Natural Resource Modeling* **4** 143–173.
- Sethi, S. P. 1979. Optimal pilfering policies for dynamic continuous thieves. *Management Sci.* **25**(6) 535–542.
- Shavell, S. 1991. Specific versus general enforcement of the law. *J. Political Econom.* **99** 1088–1108.
- Shaw, D. S., J. M. Post, K. G. Ruby. 1999. Inside the minds of the insider. *Security Management* (December) 34–44.
- Shiple, G. 1999. ISS RealSecure pushes past newer IDS players. *Network Comput.* (May 17).
- Sriram, T. 2002. Blocking virus requests in Novell bordermanager's HTTP accelerator. Feature article, Novell Appnotes, Waltham, MA.
- Stigler, G. 1970. The optimum enforcement of laws. *J. Political Econom.* **78** 526–536.
- Thomas, M. U., Y. Nisgav. 1976. An infiltration game with time dependent payoff. *Naval Res. Logist. Quart.* **23** 297–302.
- Trees, H. V. 2001. *Detection, Estimation and Modulation Theory—Part I*. John Wiley, New York.
- Tudor, J. K. 2001. *Information Security Architecture*. Auerbach Publications, Boca Raton, FL.
- Weissenberger, S. 1992. Deterrence and the design of treaty verification systems. *IEEE Trans. Systems, Man, Cybernetics* **22** 903–915.
- Yocom, B., K. Brown. 2001. Intrusion Battleground Evolves. *Network World* (October 8) 53–62.
- Zamboni, D., E. Spafford. 1999. New directions for the AAPHID architecture. *Workshop Recent Adv. Intrusion Detection*, West Lafayette, IN.

Copyright 2005, by INFORMS, all rights reserved. Copyright of Information Systems Research is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.