

# The Economics of Information Security Investment

LAWRENCE A. GORDON and MARTIN P. LOEB  
University of Maryland

---

This article presents an economic model that determines the optimal amount to invest to protect a given set of information. The model takes into account the vulnerability of the information to a security breach and the potential loss should such a breach occur. It is shown that for a given potential loss, a firm should not necessarily focus its investments on information sets with the highest vulnerability. Since extremely vulnerable information sets may be inordinately expensive to protect, a firm may be better off concentrating its efforts on information sets with midrange vulnerabilities. The analysis further suggests that to maximize the expected benefit from investment to protect information, a firm should spend only a small fraction of the expected loss due to a security breach.

Categories and Subject Descriptors: H.1.1 [Models and Principles]: Systems and Information Theory—*value of information*; K.6.0 [Management of Computing and Information Systems]: General—*economics*; K.6.5 [Management of Computing and Information Systems]: Security and Protection

General Terms: Economics, Security

Additional Key Words and Phrases: Optimal security investment

---

## 1. INTRODUCTION

Security of a computer-based information system should, by design, protect the confidentiality, integrity, and availability of the system (e.g., see NIST [1995, p. 5]). Given the information-intensive characteristics of a modern economy (e.g., the Internet and World Wide Web), it should be no surprise to learn that information security is a growing spending priority among most companies. This growth in spending is occurring in a variety of areas including software to detect viruses, firewalls, sophisticated encryption techniques, intrusion detection systems, automated data backup, and hardware devices [Larsen 1999]. The above notwithstanding, a recent study by the Computer Security Institute, with the participation of the Federal Bureau of Investigation, reported that “Ninety-one

---

This research was partially supported by The Robert H. Smith School of Business, University of Maryland and the Laboratory for Telecommunications Sciences (within the Department of Defense) through a grant with the University of Maryland Institute for Advanced Computer Studies.

Authors’ address: The Robert H. Smith School of Business, University of Maryland, College Park, College Park, MD 20742-1815; email: {lgordon; mloeb}@rhsmith.umd.edu.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2002 ACM 1094-9224/02/1100-0438 \$5.00

ACM Transactions on Information and System Security, Vol. 5, No. 4, November 2002, Pages 438–457.

percent of respondents . . . detected computer security breaches within the last twelve months” [Power 2001, p. 33]. Moreover, the study found that (for those organizations that provided loss estimates) the losses averaged over \$2 million per organization [Power 2001, p. 33]. Hence, it would appear that many firms are not adequately investing in information security. In this regard, based on a large information security survey, KPMG [2000, p. 1] concluded that, “Our overall finding is that information security requirements are not being adequately addressed, especially in the new fast moving, global, e-business environment. This will leave some organizations critically exposed.”

The importance of information security in a computer-based environment has resulted in a large stream of research that focuses on the technical defenses (e.g., encryption, access control, and firewalls) associated with protecting information (e.g., Anderson [1972], Wiseman [1986], Simmons [1994], Muralidhar et al. [1995], Denning and Branstad [1996], Sandhu et al. [1996], Schneier [1996], Pfleeger [1997], Larsen [1999], Peyravian et al. [1999], and Osborn et al. [2000]) and intrusion detection systems (e.g., Denning [1987], Daniels and Spafford [1999], Vigna and Kemmerer [1999], Axelsson [2000], and Frincke [2000]). In addition, research has been rapidly developing that focuses on the behavioral aspects of reducing information security breaches (e.g., Straub [1990], Loch et al. [1992], and Straub and Welke [1998]). In contrast, research focusing on the economic aspects of information security is rather sparse. The work that does exist on, or related to economic aspects of information security provides little generic guidance on how to derive the proper amount to invest on such security (e.g., see Millen [1992], Luotonen [1993], McKnight et al. [1997], Finne [1998], Jones [1997], Buzzard [1999], Hoo [2000], Anderson [2001], Meadows [2001], and Powers [2001]).

The purpose of this article is to derive an economic model that determines the optimal amount to invest in information security. Accordingly, information security in our model may be broadly interpreted. Our model is applicable to investments related to various information security goals, such as protecting the confidentiality, availability, authenticity, non-repudiation, and integrity of information.<sup>1</sup> Although there is often a conflict among these goals, the model we present does not address this conflict. Rather, we construct a model that specifically considers how the vulnerability of information and the potential loss from such vulnerability affect the optimal amount of resources that should be devoted to securing that information. Without a careful analysis of the effect of vulnerability on information security, intuition might suggest that, for a given potential loss and a given threat level, the optimal amount to spend on such security is an increasing function of the information’s vulnerability. Our analysis demonstrates that this may, or may not, be the case.<sup>2</sup> We demonstrate

<sup>1</sup>Moreover, our model could be used to gain insights for the optimal protection of assets other than information.

<sup>2</sup>This is in contrast with earlier literature, such as Pfleeger [1997, Chapter 10], which discuss the importance of vulnerability in the decision to invest in information security, but does not examine the effects of changes in vulnerability on the optimal investment in information security. Previous papers on information security usually combine vulnerability with the potential dollar loss associated with such vulnerability, to come up with the notion of risk (e.g., Straub and Welke [1998] and

that under certain sets of assumptions concerning the relationship between vulnerability and the marginal productivity of the security investment, the optimal investment in information security may either be strictly increasing or first increase and then decrease as vulnerability increases. Thus, under plausible assumptions, investment in information security may well be justified only for a midrange of information vulnerabilities. That is, little or no information security is economically justified for extremely high, as well as extremely low, levels of vulnerability.

Our analysis also indicates that, even within the range of justifiable investments in information security, the maximum amount a risk-neutral<sup>3</sup> firm should spend is only a fraction of the expected loss due to security breaches. For two broad classes of security breach probability functions, this fraction never exceeds 37% of the expected loss. For most cases, however, this fraction is substantially below the 37% level. Given that organizations possess limited resources, our analysis provides managers with a framework for considering decisions regarding the allocation of scarce information security dollars.

The remainder of the article is organized as follows: In the second section, vulnerability is formally defined and the general model is presented. The third section contains an analysis of how vulnerability affects the optimal level of investment in information security, given the potential loss associated with such vulnerability. The fourth, and final, section of the article offers some concluding comments.

## 2. THE MODEL

We consider a one-period model<sup>4</sup> of a firm contemplating the provision of additional security to protect a given information set. The information set could take many forms, such as a list of customers, an accounts payable ledger, a strategic plan, or company website. The increased security could be with respect to protecting the confidentiality, integrity, authenticity, non-repudiation, or availability to authorized users of the information set. An information set is characterized by three parameters:  $\lambda$ ,  $t$ , and  $v$ , representing, respectively, the loss conditioned on a breach occurring, the probability of a threat occurring, and the vulnerability, defined in the model as the probability that a threat once realized (i.e., an attack) would be successful.

---

Finne [1998]). Thus, earlier literature entangles the relationship between information vulnerability and the proper amount to spend on preventing such vulnerability.

<sup>3</sup>If someone is **risk-neutral**, it means that they are indifferent to investments that have the same expected value, even though the investments may have varying amounts of risk. Thus, a risk-neutral decision-maker would be indifferent to Investment #1 that generates either a net return of \$200,000 or a net loss of \$100,000 each with probability of 0.5, and Investment #2 that generates a net return of either \$40,000 or \$60,000 each with probability of 0.5, as both investments have an expected net return of \$50,000. Notice that Investment #1 has more risk (i.e., larger standard deviation around the expected value) than investment #2, and yet the two investments are being considered equal. Someone who is **risk-averse** would require a higher expected value for an investment with a higher risk.

<sup>4</sup>In one-period economic models, all decisions and outcomes occur in a simultaneous instant. Thus, dynamic aspects, such as a first-mover advantage or the time value of money, are not considered.

ACM Transactions on Information and System Security, Vol. 5, No. 4, November 2002.

The parameter  $\lambda$  represents the monetary loss to the firm caused by a breach of security of the information set. This loss could be due to a security breach related to confidentiality (e.g., the loss due to the strategic information becoming available to competitors or the fraudulent use of credit card information by hackers), integrity (e.g., the loss due to the firm making faulty decisions based on data altered by an intruder), or denial of services (e.g., loss due to missed sales from authorized users who were denied legitimate access). Although  $\lambda$  would normally depend on the use of the information (by the firm itself, by competitors, or by hackers) and would change over time, for simplicity we take  $\lambda$  to be a fixed amount as estimated by the firm (e.g., the present value of lost profits from current and future lost sales). Even though we initially assume that this loss is a fixed value, we will investigate how changes in the value of the loss affect the firm's security investment decision. However, we assume  $\lambda$  is finite and less than some very large number, say  $M$ .<sup>5</sup> Thus, the model is not intended to cover protection of national/public assets or other circumstances where a loss could be catastrophic.

The probability of an attempted breach of the given information set is denoted by  $t \in [0, 1]$ , and we call  $t$  the threat probability. We make the simplifying assumption that there is a single threat to an information set.<sup>6</sup> The parameter  $v$  is used to denote the information set's vulnerability, by which we mean the probability that *without* additional security, a threat that is realized will result in the information set being breached and the loss,  $\lambda$ , occurring. Our view of threats and vulnerabilities is consistent with the argument of Littlewood et al. [1993, p. 228] concerning "the desirability of a probability-based framework for operational security measurement." Since  $v$  is a probability,  $v \in [0, 1]$ .

Typically, the threat to an information set and the information set's vulnerability would lie in the interior (i.e.,  $0 < t < 1$  and  $0 < v < 1$ ). Note that the information is completely invulnerable when  $v = 0$ . One can consider an information set on a computer buried in concrete thirty feet underground to be completely invulnerable. Of course, this state of invulnerability (and perfect confidentiality) is achieved at the cost of having the information set become completely inaccessible.<sup>7</sup> Similarly, if  $v = 1$ , the information set is completely vulnerable. Such information sets, like last quarter's statement of earnings (for a publicly traded firm) or the retail price of a specific product, may be viewed as public information. For a given information set, the probability of the loss occurring (sometimes called the risk of the loss) is the product of the vulnerability and the threat probabilities. Thus, the product  $vt\lambda$  represents the expected

<sup>5</sup>For a catastrophic loss,  $\lambda \geq M$ , the assumption of risk-neutrality becomes unrealistic. In the language of economics, the disutility of a catastrophic loss is so large that decision makers would prefer the expected value of the gamble rather than risking a loss of  $\lambda$ .

<sup>6</sup>Allowing multiple threats significantly increases the complexity of the model. However, there is no reason to believe that a more complex economic model would yield additional insights. In fact, it is often argued that clearer insights are provided by models that are less rather than more complex. In this vein, Varian [1997, p. 4] writes, "A model is supposed to reveal the essence of what is going on: your model should be reduced to just those pieces that are required to make it work."

<sup>7</sup>Hence, this is one illustration of the trade-offs among the goals of confidentiality, integrity, and availability referred to earlier.

loss (conditioned on no investment in information security) associated with the given information set.<sup>8</sup> Thus, for any positive threat ( $t > 0$ ), the expected loss increases with the vulnerability.

Of course, firms can and do invest in information security.<sup>9</sup> In general, one would expect a firm to have more influence over an information set's vulnerability than over the threats to the information set.<sup>10</sup> For the purposes of our model, we make the simplifying assumption that firms can influence the vulnerability of an information set by investing in information security, but the firm cannot invest to reduce the threat. We therefore fix the threat probability at  $t > 0$ , and focus on the firm's choice of the level of investment to reduce the vulnerability of their information.<sup>11</sup> Since the threat probability is held constant, for notational simplicity we define  $L = t\lambda$ . For expositional ease, we will refer to  $L$  as the loss or potential loss associated with the information set.

Let  $z > 0$  denote the monetary (e.g., dollar) investment in security to protect the given information set. Thus,  $z$  is measured in the same units (i.e., dollars) used to measure the potential loss  $L$ . The purpose of the investment  $z$  is to lower the probability that the information set will be breached. Let  $S(z, v)$  denote the probability that an information set with vulnerability  $v$  will be breached, conditional on the realization of a threat and given that the firm has made an information security investment of  $z$  to protect that information. We refer to the function  $S(z, v)$  as the security breach probability function and to its value at a particular level of  $z$  and  $v$  as the security breach probability.

As is common with nearly all economic models, we abstract from reality and assume that postulated functions are sufficiently smooth and well behaved. This is done so that an optimization problem, which can be solved with basic tools of calculus, can be used to represent the economic phenomenon. In our model, we assume that the function  $S(z, v)$  is continuously twice differentiable. Of course, in reality, discrete investments in new security technologies are often necessary to get any incremental result. Such discrete investments result in discontinuities. However, even though the commitment to invest in security may be made in discrete pieces, the actual expenditures can often be broken down into small increments. Furthermore, some information investments can be reversed (e.g., additional security personnel can be fired and purchased equipment and software can be sold). Thus, a smooth approximation of the security investment

<sup>8</sup>As noted in the previous footnote, the calculation of the expected loss becomes more complicated when multiple threats are considered. Assume for simplicity that a threat that results in a breach causes a loss of  $\lambda$ , but that there can be no additional losses from a second breach (once you're shot dead, additional threats are irrelevant). Now suppose there are two (independent) threats occurring with probability  $t_1 = 0.8$  and  $t_2 = 0.9$  and suppose the vulnerability probability is  $v = 0.1$ . Then, the probability of a loss (calculated using a simple decision tree) will be  $0.1628 < vt_1 + vt_2$ .

<sup>9</sup>Investments in information security have many of the same characteristics of what firms usually consider capital expenditures. This fact notwithstanding, firms usually treat an inordinate portion of the costs of information security as operating expenditures. Although beyond the scope of this paper, such treatment raises its own set of interesting questions.

<sup>10</sup>Of course, this may not always be the case. For example, if each employee having access to an information set is viewed as a threat, the threat can be reduced by restricting employee access.

<sup>11</sup>Although we hold  $t$  fixed, our model allows us to see how changes in the value of the parameter  $t$  (and the parameter  $\lambda$ ) would change the optimal security investment decision.

represents a reasonable first approach to gaining insights into the problem of determining the optimal investment in information security.<sup>12</sup>

The nature of information vulnerability and information security leads us to consider the following assumptions concerning  $S(z, v)$ :

**A1.**  $S(z, 0) = 0$  for all  $z$ . That is, if the information set is completely invulnerable, then it will remain perfectly protected for any amount of information security investment, including a zero investment.

**A2.** For all  $v$ ,  $S(0, v) = v$ . That is, if there is no investment in information security, the probability of a security breach, conditioned on the realization of a threat, is the information set's inherent vulnerability,  $v$ .

**A3.** For all  $v \in (0, 1)$ , and all  $z$ ,  $S_z(z, v) < 0$  and  $S_{zz}(z, v) > 0$ , where  $S_z$  denotes the partial derivative with respect to  $z$  and  $S_{zz}$  denotes the partial derivative of  $S_z$  with respect to  $z$ . That is, as the investment in security increases, the information is made more secure, but at a decreasing rate. Furthermore, we assume that for all  $v \in (0, 1)$ ,  $\lim_{z \rightarrow \infty} S(z, v) = 0$ , so by investing sufficiently in security, the probability of a security breach,  $t$  times  $S(z, v)$ , can be made to be arbitrarily close to zero.

Note that from A3 that even a very small expenditure for information security will reduce the probability of a security breach. This may be due to the fact that there are no fixed costs of information security. An alternative interpretation of the model views the investment in information security as an incremental investment beyond security measures already in place. A firm may have an Information Technology Director and other IT staff who devote limited time to security issues. By allocating a bit more time (and hence money) to security issues, it would be reasonable to expect some decrease in the probability of a breach. Similarly, most firms have some security measures (e.g., firewalls, intrusion detection systems, antivirus software) in place and are considering incremental expenditures to enhance or supplement these measures. Also, note that A3 implies that no finite investment in information security can make a vulnerable ( $v > 0$ ) information set perfectly secure. The analysis that follows assumes that the security breach probability functions meet assumptions A1–A3.

In order to determine the amount to invest in information security, a risk-neutral firm would compare the expected benefits of the investment with cost of the investment.<sup>13</sup> The expected benefits of an investment in information

<sup>12</sup>By making such simplifying assumptions, economists have been able to gain powerful insights that have proven valid in more general settings.

<sup>13</sup>We believe risk-neutrality is a reasonable assumption for most security-related issues. Of course, if the loss associated with a security breach were of an immense magnitude, a more realistic assumption may well be that of risk-aversion. By implicitly restricting the magnitude of the potential loss, we concur with Littlewood et al. [1993, p. 217], who write, "in these initial stages of attempting to model operational security, we should restrict ourselves to systems for which the security requirements are also modest." Under a risk-averse assumption, the level of expenditure on information security would depend on the specific nature and degree of the decision-maker's risk aversion (modeled by economists as the decision-maker's utility function), and the optimal investment in information security would increase with the level of risk-aversion. Such an analysis, however, is beyond the scope of this article.



security, denoted as *EBIS*, are equal to the reduction in the firm's expected loss attributable to the extra security. That is:

$$EBIS(z) = [v - S(z, v)] L. \quad (1)$$

*EBIS* is written above as a function of  $z$ , since the investment in information security is the firm's only decision variable ( $v$  and  $L$  are parameters of the information set). The expected net benefits from an investment in information security, denoted *ENBIS* equal *EBIS* less the cost of the investment, or:

$$ENBIS(z) = [v - S(z, v)] L - z. \quad (2)$$

To focus on the effect of vulnerability, we denote the optimal investment as  $z^*(v)$ . Observe that from A1, if an information set is completely invulnerable, the optimal investment in information security is set equal to zero, that is,  $z^*(0) = 0$ . For now, we assume that the information set is neither completely vulnerable nor completely invulnerable, that is,  $0 < v < 1$ .

From Assumption (A3),  $S(z, v)$  is strictly convex in  $z$ , thus *ENBIS* is strictly concave in  $z$ . Hence, an interior maximum  $z^* > 0$  is characterized by the first-order condition:

$$-S_z(z^*, v) L = 1. \quad (3)$$

where the left hand side of (3) represents the marginal benefits from the security investment and the right hand side of (3) represents the marginal cost of investment.<sup>14</sup> One should invest in security only up to the point where marginal benefit equals marginal cost.

Recall that the value of an information set is measured by the potential loss associated with the information set. It follows from Eq. (3), as one would expect, that for a given level of vulnerability, the optimal amount to be invested in information security,  $z^*$ , increases with increases in the value of the information set (i.e., with increases in the threat  $t$  or the loss  $\lambda$ ).<sup>15</sup>

This optimal level of investment in information security is illustrated in Figure 1. From Eq. (1), A1, and A2, the benefits of an investment in information security, *EBIS*( $z$ ), start out at zero and approach  $vL$  as the investment level increases. The costs of the investment are given by  $z$ , the 45° line in Figure 1.

<sup>14</sup>Recall that  $z$  measures information security investment in dollars (or other monetary units). Hence, by definition, the price of a unit of  $z$  equals one. Thus, the marginal cost of investment (i.e., the cost of increasing  $z$  by one unit) equals one.

<sup>15</sup>This can be seen by first rewriting (3) as:

$$-S_z(z^*, v) = \frac{1}{L}$$

and taking the total differential to get:

$$S_{zz}(z^*, v) dz^* = \frac{dL}{L^2}.$$

This yields:

$$\frac{dz^*}{dL} = \frac{1}{L^2 S_{zz}(z^*, v)}.$$

Thus, as  $S_{zz}(z^*, v)$  is positive from assumption A3, we have  $(dz^*/dL) > 0$ , giving the desired result.

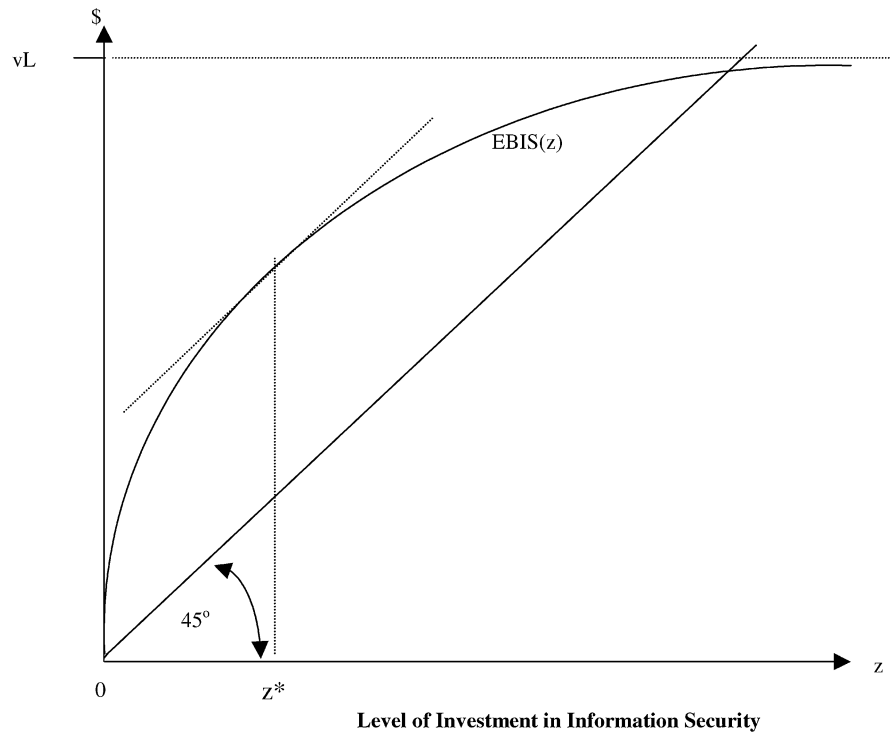


Fig. 1. The benefits and cost of investment in information security.

The optimal investment,  $z^*$ , is where the difference between benefits and costs are maximized, and at that point the tangent to  $EBIS(z^*)$ , has a slope, representing the marginal benefits, equal to the marginal cost of one. Observe that the optimal amount to be invested in information security,  $z^*$ , is less than  $vL$ , the loss that would be expected in the absence of any investment in security.<sup>16</sup> This can be seen by noting in Eq. (2) that the expected benefits will always be less than  $vL$ . In Figure 1, this can be seen by noting that the benefits of an investment in information security,  $EBIS(z)$ , crosses the  $45^\circ$  line below  $vL$ . In section 3, for two broad classes of security breach probability functions, we show that the optimal amount to be invested in information security is only a small fraction of the expected loss,  $vL$ .

The optimal level of investment in information security equals zero if the marginal benefits at  $z = 0$  are less than or equal to the marginal costs of such investment. This condition can be rewritten as:

$$L \leq \frac{1}{-S_z(0, v)}. \tag{4}$$

Since our focus is on the effects of vulnerability, we are interested in determining the levels of  $v$  that cause the optimal level of investment in information

<sup>16</sup>To see this formally, note that  $0 < vL - S(z^*, v)L - z^* < vL - z^*$ , so  $z^* < vL$ .



security to become zero, holding  $L$  constant.<sup>17</sup> For a given  $L$ ,  $z^*(v) = 0$  whenever  $-S_z(0, v)$ , a positive number, is sufficiently small.

### 3. HOW VULNERABILITY AFFECTS THE OPTIMAL LEVEL OF INVESTMENT IN INFORMATION SECURITY

We now investigate the properties of  $z^*(v)$  to see how vulnerability affects the optimal level of investment in information security. From the first-order condition given in Eq. (3), we see that vulnerability affects the optimal level of investment by affecting the partial derivative of the security breach function with respect to  $z$ . This partial derivative,  $S_z(z, v)$ , may be interpreted as the marginal productivity of security investment, as it measures the rate at which the probability of a security breach decreases with an increase in security investment. Thus, the change in the optimal level of information security investment in response to a change in vulnerability is determined by the cross partial derivative  $S_{zv}(z, v)$ , which may be interpreted as the change in the marginal productivity of the investment with respect to a change in vulnerability.

If the information set were perfectly invulnerable ( $v = 0$ ), then no investment in information security would be made (i.e.,  $z^*(0) = 0$ ). At some sufficiently larger level of vulnerability, it would be optimal to make a positive investment in information security in order to reduce the probability of the loss (and, therefore the expected loss). Hence, in some range, an increase in vulnerability leads to an increase in investment in information security. This observation is stated in the following proposition (a formal proof appears in the appendix).

**PROPOSITION 1.** *For all security breach probability functions for which A1–A3 hold, there exists a loss,  $L$ , and a range of  $v$  in which increases in vulnerability result in an increase in the optimal investment in information security.*

In order to be able to calculate a closed form solution for  $z^*(v)$  and gain further insights into the relationship between vulnerability and optimal security investment, we examine two broad classes of security breach probability functions. The first class of security breach probability functions, denoted by  $S^I(z, v)$ , is given by:

$$S^I(z, v) = \frac{v}{(\alpha z + 1)^\beta} \quad (5)$$

where the parameters  $\alpha > 0$ ,  $\beta \geq 1$  are measures of the productivity of information security (i.e., for a given  $(v, z)$ , the probability of a security breach is decreasing in both  $\alpha$  and  $\beta$ ). As is easily verified, each member of this class of security breach probability functions satisfies conditions A1–A3, and was selected because of its relatively simple functional form. In particular, the security breach probability functions in this class is linear in vulnerability. Figure 2 shows how increases in the amount of investment in information security,  $z$ , reduce the expected loss from an information security breach. The top

<sup>17</sup>Clearly, if one were to hold  $v$  constant and let  $L$  vary, the optimal investment in information security will be zero for sufficiently small  $L$ . That is, if the loss conditional on a security breach is very small, a positive investment in information security is not justified.

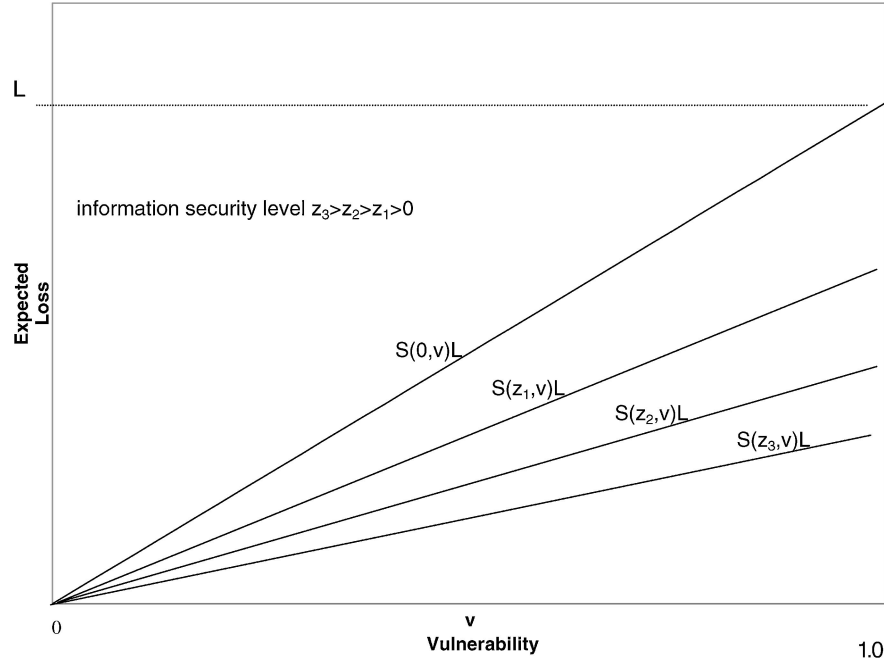


Fig. 2. Expected value of information loss,  $S(z, v)L$ , as vulnerability increases at different levels of investment in information security (for Class I).

line,  $S(0, v)L$  in Figure 2, equals  $vL$ , the expected loss without increased investment in information security. The straight line below it represents  $S(z_1, v)L$ , which is the expected loss when  $z_1$  is invested in information security. Thus, for an information set with vulnerability  $v$ , the difference between the lines at  $v$  represents *EBIS* (i.e., the expected benefit of investing  $z_1$  in information security gross of the costs of the investment).

For security breach probability functions belonging to this first class, an expression for an interior optimal level of investment in information security can be found by solving for  $z^*$  in the first-order condition given by Eq. (3). Letting  $z^{I^*}(v)$  denote this optimal yields:

$$z^{I^*}(v) = \frac{(v\beta\alpha L)^{1/(\beta+1)} - 1}{\alpha}. \quad (6)$$

For this first class of security breach probability functions, condition (4) yields that  $z^{I^*}(v) = 0$  for  $0 \leq v \leq 1/\alpha\beta L$ . Thus, for the first class of security breach functions, the optimal investment in security equals zero until  $v = 1/\alpha\beta L$ , and then, based on Eq. (6), increases at a decreasing rate (see Figure 3). As  $z^{I^*}(v)$  is strictly increasing in  $v$  over the high range of vulnerabilities, Figure 3 illustrates that, at least for security breach probability functions belonging to  $S^I(z, v)$ , for a given potential loss, a firm can be better off concentrating its resources on high-vulnerability information sets.

We now examine a second broad class of security breach probability functions that also meets assumptions A1–A3, yet demonstrates that a firm is not




Fig. 3. Optimal value of security investments as a function of vulnerability,  $z^*(v)$  for Class I.

always better off concentrating its resources on high vulnerability information sets. Consider the second class of security breach probability functions is given by:

$$S^{II}(z, v) = v^{\alpha z + 1} \quad (7)$$

where the parameter  $\alpha > 0$ , is a measure of the productivity of information security. Each curved lines in Figure 4 represents a particular member of the class  $S^{II}(z, v)$ , parameterized by varying values of  $z > 0$ , for a fixed level of  $\alpha$ . At any level of vulnerability,  $v$ , the difference between one of the curved lines and the straight line (representing  $vL$ ) gives the *EBIS* for the given investment,  $z$ , in securing the confidentiality of the information set. This class of security breach probability functions has the property that the cost of protecting highly vulnerable information sets becomes extremely expensive as the vulnerability of the information set becomes very large.<sup>18</sup>

Using the first-order condition given in Eq. (3), the expression for the interior optimal level of investment in information security for  $S^{II}(z, v)$  is found to be:

$$z^{II*}(v) = \frac{\ln(1/\alpha v L(\ln v))}{\alpha \ln v} \quad (8)$$

<sup>18</sup>The class of security breach function  $S^{II}(z, v)$  given in Eq. (7) is not the only class of security breach functions that has this property and could be used to demonstrate the propositions that are given later in this section. For example, the class of security breach probability functions given by  $S^{III}(z, v) = v e^{\alpha z(v-1)}$ , where  $\alpha > 0$  could have been used instead of  $S^{II}(z, v)$ . The class  $S^{II}(z, v)$  was selected for presentation because of its slightly simpler form.

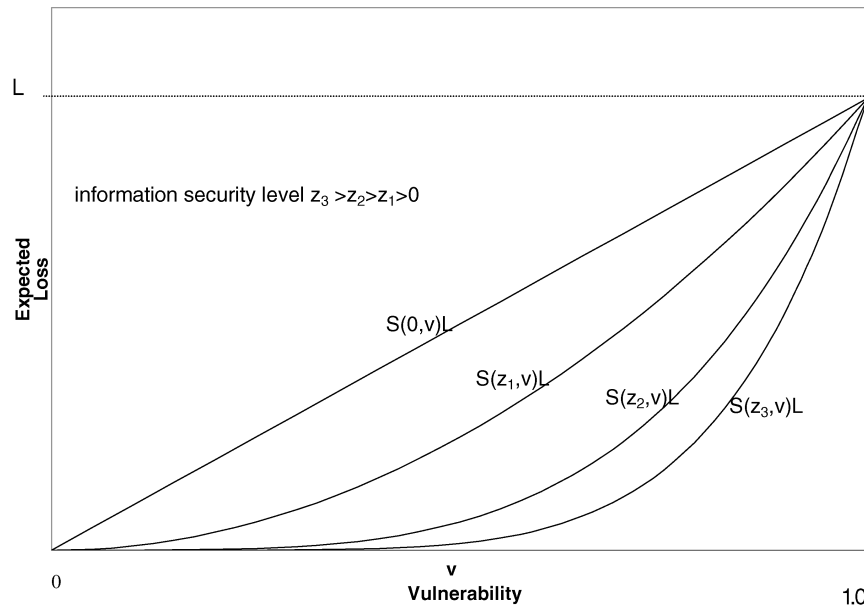


Fig. 4. Expected value of information loss,  $S(z, v)L$ , as vulnerability increases at different levels of investment in information security (for Class II).

For this second class of security breach probability functions, condition (4) can be rewritten (after rearranging terms) as  $1/L > -\alpha v \ln v$ . Note that  $-\alpha v \ln v > 0$  for  $0 < v < 1$ , and takes on a maximum at  $v = 1/e \approx 0.3679$ , and gets sufficiently close to 0 for  $v$  sufficiently close to either 0 or 1. Thus, for a given  $L$ , there exists a lower limit,  $\underline{V}(L)$ , and an upper limit  $\overline{V}(L)$ , with  $0 < \underline{V}(L) < \overline{V}(L) < 1$ , such that  $z^{II*}(v) = 0$ , when  $0 < v < \underline{V}(L)$  or  $\overline{V}(L) < v < 1$  and  $z^{II*}(v) > 0$  when  $\underline{V}(L) < v < \overline{V}(L)$ . Although one cannot find a closed form expression for  $\underline{V}(L)$  and  $\overline{V}(L)$ , by plotting  $z^{II*}(v)$ , numerical values for these points can easily be approximated.<sup>19</sup> The regions of extremely low and extremely high vulnerability are shown in the graph of  $z^*(v)$  for  $S^{II}(z, v) = v^{\alpha z + 1}$  in Figure 5.

While our earlier proposition (and the analysis of the first class of security breach probability functions) left open the possibility that the optimal investment in information security is always (weakly) increasing in vulnerability, the analysis of the second class of security breach probability functions shows that this is not the case. We have seen that the class of security breach probability functions  $S^{II}(z, v) = v^{\alpha z + 1}$  meets conditions A1–A3 and results in the optimal security investment first increasing and then decreasing in the vulnerability. Thus, the demonstration and analysis of the second class of security breach probability functions provides a counterexample that is sufficient to prove the following:

**PROPOSITION 2.** *Suppose a security breach probability function meets conditions A1–A3, then it is not necessarily the case that the optimal level of investment in information security,  $z^*(v)$ , is weakly increasing in vulnerability,  $v$ .*

<sup>19</sup>For example, when  $\alpha = 0.00001$ , and  $L = \$400,000$ , then  $\underline{V} \approx 0.1$  and  $\overline{V} \approx 0.7$ .




Fig. 5. Optimal value of security investments as a function of vulnerability,  $z^*(v)$  for Class II.

Proposition 2 indicates that a firm should be careful in deciding where to concentrate information security resources. Figures 3 and 5 illustrate that for a given potential loss, a firm may be better off concentrating its resources on high vulnerability information sets (as demonstrated by the fact that for the first class of security breach probability functions,  $z^{I*}(v)$  is strictly increasing in  $v$  over the high range of vulnerabilities), or on information sets with midrange vulnerabilities (as demonstrated by  $z^{II*}(v)$  for  $S^{II}(z, v)$ ). In other words, for the second class of security breach probability functions (which meets assumption (A4)), the area of zero investment, for a given  $L$ , should be two tailed rather than one tailed. For security breach functions in class II, as well as in class I, the marginal benefit from investment in information security for low vulnerability information sets does not justify the investment since the security of such information is already good. For security breach functions in class II, when an information set is extremely vulnerable, the benefit of spending a given amount for increased information security of the information (as measured by the decrease in expected loss from the extra security) is very small. For example, for the case where the security issue is that of confidentiality, knowledge that a firm is trying to sell a particular business unit may become nearly public information. In such a case, because of the multiple sources of potential information leakage, it may well be too expensive to monitor employees and business contacts to provide even a mild level of information security. Hence, the key in analyzing information security decisions is not the vulnerability (or the expected loss without the investment), but the reduction in expected loss with the investment.

The next proposition provides insight into the relationship between the optimal level of investment in security and the loss that would be expected in the absence of any investment in security when the security probability breach functions belong to class I or class II.

**PROPOSITION 3.** *Suppose the security breach probability function belongs to class I (i.e., it can be expressed as  $S^I(z, v) = v/(\alpha z + 1)^\beta$  for some  $\alpha > 0$ ,  $\beta \geq 1$ ) or to class II (i.e., it can be expressed as  $S^{II}(z, v) = v^{\alpha z + 1}$  for some  $\alpha > 0$ ), then  $z^*(v) < (1/e)vL$ . (See Appendix for proof.)*

Proposition 3 shows that, for the two broad classes of information security breach probability functions, the optimal investment in information security is always less than or equal to 36.79% of the loss that would be expected in the absence of any investment in security.<sup>20</sup> The restriction that the security breach probability functions have one of two specific functional forms warrants a discussion of the robustness of the proposition. First, note that the two classes of security breach probability functions appear unrelated, other than the fact that functions in both classes satisfy conditions A1–A3. Functions belonging to class I are linear in vulnerability and those belonging to class II are strictly concave (for  $\alpha > 0$ ). Moreover, the result holds for all values of  $\alpha > 0$ ,  $\beta \geq 1$ , that is, the productivity of information security is unrestricted.<sup>21</sup> Second, the proposition critically depends on the assumption that the firm already has some information security infrastructure in place (e.g., an IT officer devoting some time to security issues, access controls, etc.) so that there are no incremental fixed costs associated with new security investments.<sup>22</sup>

The practical import of Proposition 3 as guidance for decision-making is enhanced when one considers that the 36.79% figure is a maximum, and for a wide range of security breach probability functions belonging to class I and II, the optimal amount to be invested in information security is considerably less. For example, for class I security breach probability functions with  $\beta = 1$ , the maximum percent to be invested is 25% of  $vL$  (as can be seen by examining Eq. (A4) in the Appendix) and only occurs when  $\alpha vL = 4$ . Thus, when  $\beta = 1$ ,  $\alpha = 0.00001$ ,  $L = \$400,000$  and  $v = 1$ , the 25% limit will hold, but at lower values of  $v$ , the optimal level of investment is less than the 25% of  $vL$ .<sup>23</sup>

The findings discussed in this section of the article can be summarized as follows: The optimal expenditures for protecting a given information set do

<sup>20</sup>As indicated in footnote 18 above, Proposition 3 extends beyond the two classes of information security breach functions.

<sup>21</sup>Also note that some simple perturbations of the two classes of security probability functions do not affect the conclusion of Proposition 3. Specifically, suppose A3 is generalized so that  $v - w$  of the probability of breach is due to sources that cannot be reduced through investment in security, that is,  $\lim S(z, v) \rightarrow v - w$ , as  $z \rightarrow \infty$ , where  $0 \leq w \leq v$ . Letting  $S(z, v) = v - w + S^i(z, w)$ , for  $i = I$ , one can easily verify that the conclusion of proposition still holds.

<sup>22</sup>Of course, if there were incremental fixed costs of  $F$ , in addition to variable costs  $z$ , then (for the two classes of breach functions), the optimal total amount spent on information security as a fraction of the expected loss in the absence of additional security would increase by  $F/vL$ . As  $F$  increases, the lower range of vulnerabilities in which investment is uneconomical increases. Clearly, if  $F$  were sufficiently large, no investment would take place.

<sup>23</sup>For example, when  $v = 0.52$ , the optimal investment is \$41,421 or 20.7% of  $vL$ .



not always increase with increases in the information set's vulnerability. Furthermore, for two broad classes of security breach probability functions, the optimal amount to invest in information security should not exceed 37% ( $\approx 1/e$ ) of the expected loss due to a security breach. The analysis presented is not without limitations. First, our result giving the maximum amount of the optimal investment in information security depended on the specific functional forms of the security breach functions and assumed no lumpiness in expenditures for information security. While the assumption that incremental fixed costs of information security investment is zero clearly played a crucial role in our demonstration, it is an open question as to whether or not our result extends to all continuous security breach functions meeting assumptions A1–A3. Second, there is no simple procedure to determine the probabilities of the threat and the vulnerability associated with an information set. Third, in a similar vein, procedures for deriving and considering the potential loss from an information security breach, especially for a huge loss (as would likely be the case for the protection of many national/public assets), is also problematic. A fourth limitation of this research is that we have not modeled how conflicts of interest between senior management and the firm's chief information security officer would affect the derivation of the optimal amount to invest in information security.<sup>24</sup> Finally, we have not modeled the case where a single investment in information security is used to protect the security of multiple information sets having correlated security risks.<sup>25</sup>

#### 4. CONCLUDING COMMENTS

The new computer-based information age has changed the way organizations operate, as well as the way they need to look at information security. Indeed, information security has become at least as important to modern corporations as is the protection of tangible physical assets. Not surprisingly, a rapidly growing body of research addresses the issue of information security. This research has focused primarily on the technical aspects of protecting information in a computer-based system (i.e., encryption, data and software controls, and hardware controls). The behavioral aspects of preventing information security breaches have also been attracting much recent attention among researchers. In contrast, very little work has been done which addresses the economic aspects of information security. In particular, given the amount of resources

<sup>24</sup>In another context, Hann and Weber [1996] model the conflict of interest between senior management and the CIO. The cost of the conflict of interest between a principal (e.g., a senior manager) and an agent (e.g., the CIO) is known in economics as an agency cost. Agency costs arise in a variety of other situations where the decision making authority is delegated by a principal (e.g., an owner) to an agent (e.g., a senior manager).

<sup>25</sup>Similarly, our article does not address the joint protection of information sets along with tangible assets such as desks, printers, and personnel. For example, fire protection adds to the security of non-information assets along with information assets. Of course, if we bundle all assets together as a single set, we could still use our model for guidance in determining a joint level of (information plus noninformation) security investment. However, our model does not give guidance on how the total investment in security should be allocated between information security investments and security investments for other assets.

currently being devoted by organizations to shore up information security, what is needed is a conceptual framework to help derive an optimal level of information security spending. This article helps to fill this void in the literature by presenting such a framework, in the form of an economic model for information security investment decisions. An economics perspective naturally recognizes that while some investment in information security is good, more security is not always worth the cost. The model given in this article specifically considers how the vulnerability of information, and the loss associated with such vulnerability, affect the optimal level of resources that should be devoted to securing information.

The analysis contained in this article has shown that, for a broad class of security breach probability functions, the optimal amount to spend on information security is an increasing function of the level of vulnerability of such information. Our analysis also shows that, for a second broad class of security breach probability functions, the optimal amount to spend on information security does not always increase with the level of vulnerability of such information. For this second class, the optimal amount to spend on information security initially increases, but ultimately decreases with the level of vulnerability of such information. Thus, the second class of security breach probability functions also shows that managers allocating an information security budget should normally focus on information that falls into the midrange of vulnerability to security breaches. Hence, a meaningful endeavor for managers may be to partition information sets into low, middle, and high levels of security breach vulnerability. Some information sets may be so difficult to protect to a very high level of security, that one may be best off defending them only at a moderate level.

Information security vendors and consultants will naturally focus on huge potential losses from security breaches in order to sell their products and services. Astute information security managers no doubt are aware that expected losses are typically an order of magnitude smaller than such potential losses. Our analysis shows that for two broad classes of security breach probability functions, the optimal amount to spend on information security never exceeds 37% of the expected loss resulting from a security breach (and is typically much less than 37%). Hence, the optimal amount to spend on information security would typically be far less than even the expected loss from a security breach.

Our findings for the two classes of security breach probability functions shed significant light on the much overlooked issue of determining how much to invest in information security. While our analysis provides new insights, a number of important aspects of the information security investment decisions are not addressed by our model, and therefore represent opportunities for extending the line of research pursued in this article. One aspect that our model does not address is the various perverse economic incentives (e.g., externalities arising when decisions of one party affects those of others) affecting investment in information security. The nature and effects of perverse economic incentives is the principle focus of a stimulating paper by Anderson [2001], and it would be interesting to examine how these incentives affect the analysis resulting from our model. As a model of a single-decision maker, our analysis does not take into

account how potential attackers of an information system change strategies in reaction to an additional security investments. That is, our analysis does not consider the game theoretic aspects of information security, although such consideration would enrich our analysis.<sup>26</sup> While our single-period model allows us to see the effects of changes in the model's parameters (e.g., the loss associated with a security breach), it would be interesting to extend our model to include dynamic issues.

In addition to extending our model as suggested above, future research could, and should, empirically assess whether or not organizations invest in information security in a manner that is consistent with the findings of this article. Of course, the differences between the empirical evidence and the analytical findings of this article would need to be explained. In this regard, particular attention should be given to determining how firms estimate the potential loss and the probabilities associated with the threats and vulnerabilities of information. The above notwithstanding, the analysis contained in this article provides a framework for future research addressing issues related to the economics of investment in information security.

## APPENDIX

**PROOF OF PROPOSITION 1.** Observe from (A1),  $S_z(z, 0) = 0$  for all  $z > 0$  and from (A3),  $S_z(z, v) < 0$ , for all  $z > 0$  and  $0 < v < 1$ . Therefore, at least over some range,  $S_z(z, v)$  is decreasing in  $v$ . Consider the pair  $(\underline{z}, \underline{v})$ , which is in the range where  $S_z(z, v)$  is decreasing in  $v$ . There exists an  $L$  such that  $-S_z(\underline{z}, \underline{v})L = 1$ , so for that  $L$ ,  $z^*(\underline{v}) = \underline{z}$ . Thus, for sufficiently small but positive  $\varepsilon$ ,  $-S_z(z^*(\underline{v}), \underline{v} + \varepsilon)L > 1$ . From (A3),  $S_{zz} > 0$ , so there exists  $\delta > 0$  such that  $-S_z(z^*(\underline{v}) + \delta, \underline{v} + \varepsilon)L = 1$ , that is,  $z^*(\underline{v} + \varepsilon) = z^*(\underline{v}) + \delta$ . Hence,  $z^*$  is increasing at  $\underline{v}$ .  $\square$

**PROOF OF PROPOSITION 3.** Suppose the security breach probability function belongs to class I. Then, using Eq. (6), we have:

$$\frac{z^{I^*(v)}}{vL} = \frac{(\beta\alpha vL)^{1/(\beta+1)} - 1}{\alpha vL}. \quad (\text{A.1})$$

Letting  $x = \alpha vL$ , Eq. (A.1) can be rewritten as:

$$\frac{z^{I^*(v)}}{vL} = \frac{(\beta x)^{1/(\beta+1)} - 1}{x}. \quad (\text{A.2})$$

The right hand side of (A.2) reaches its maximum at:

$$x = (\beta + 1)^{\beta+1} \beta^{-2-\beta}, \quad (\text{A.3})$$

and substituting this (A.3) into (A.2) we get:

$$\frac{z^*}{vL} = \left( \frac{\beta}{\beta + 1} \right)^{\beta+1}. \quad (\text{A.4})$$

<sup>26</sup>This game-theoretic aspect is noted by Jajodia and Millen [1993, p. 85], "Computer security is a kind of game between two parties, the designer of a secure system, and a potential attacker." The game-theoretic aspect of information security is also highlighted by Gordon and Loeb [2001].

The right hand side of (A.4) is increasing in  $\beta$ . Applying L'Hospital's rule, we have:

$$\lim_{\beta \rightarrow \infty} \left( \frac{\beta}{\beta + 1} \right)^{\beta+1} = \frac{1}{e}. \quad (\text{A.5})$$

Hence, the right hand side of (A.4) is less than  $1/e$  and  $z^*(v) < (1/e)vL$  for the first class of security breach probability functions.

Now suppose the security breach probability function belongs to class II. Using Eq. (8), we have:

$$\frac{z^{II^*}(v)}{vL} = \frac{\ln(1/\alpha vL \ln v)}{\alpha vL \ln v}. \quad (\text{A.6})$$

Letting  $x = -\alpha vL \ln v$ , Eq. (A.6) can be rewritten as:

$$\frac{z^{II^*}(v)}{vL} = \frac{\ln(1/x)}{-x}. \quad (\text{A.7})$$

The first-order condition for maximum of the right-hand side is:

$$\frac{1 + \ln(1/x)}{x^2} = 0. \quad (\text{A.8})$$

Condition (A.8) is satisfied at the point  $x = e$ , as is the second-order condition:

$$\frac{-3 - 2 \ln(1/x)}{x^3} < 0. \quad (\text{A.9})$$

Thus, the right-hand side of (A.7) is maximized at  $x = e$ , taking on a maximum value of  $1/e$  at this point. Hence,  $z^{II^*}(v)/vL < 1/e$ . Hence,  $z^*(v) < (1/e)vL$  also holds for the second class of security breach probability functions.  $\square$

#### ACKNOWLEDGEMENTS

The authors wish to thank Mike Ball, John Hughes, Jon Millen, Ravi Sandhu, Tashfeen Sohail, Gene Spafford, Zheng Wang and the participants at the accounting and finance workshop at the London School of Economics and Political Science for comments on an earlier version of this article.

#### REFERENCES

- ANDERSON, J. 1972. Computer security technology planning study. U.S. Air Force Electronic Systems Division Tech. Rep. (Oct.), 73–51.
- ANDERSON, R. 2001. Why information security is hard—An economic perspective. In *Proceedings of 17th Annual Computer Security Applications Conference (ACSAC)* (New Orleans, La. Dec. 10–14).
- AXELSSON, S. 2000. The base-rate fallacy and the difficulty of intrusion detection. *ACM Trans. Inf. Syst. Sec.* 3, 3 (Aug.), 186–205.
- BUZZARD, K. 1999. Computer security—What should you spend your money on. *Comput. Sec.* 18, 4, 322–334.
- DANIELS, T. E. AND SPAFFORD, E. H. 1999. Identification of host audit data to detect attacks on low-level IP. *J. Comput. Sec.* 7, 1, 3–35.

- DENNING, D. 1987. An intrusion-detection model. *IEEE Trans. Softw. Eng.* 13, 2 (Feb.), 222–226.
- DENNING, D. AND BRANSTAD, D. 1996. A taxonomy of key escrow encryption systems. *Commun. ACM*, 39, 3 (Mar.), 34–40.
- FINNE, T. 1998. A conceptual framework for information security management. *Comput. Sec.* 17, 4, 303–307.
- FRINCKE, D. 2000. Balancing cooperation and risk in intrusion detection. *ACM Trans. Inf. Syst. Sec.* 3, 1 (Feb.), 1–29.
- GORDON, L. AND LOEB, M. 2001. A framework for using information security as a response to competitor analysis systems. *Commun. ACM*, 44, 9 (Sept.), 70–75.
- HANN, J. AND WEBER, R. 1996. Information systems planning: A model and empirical tests. *Management Sci.* 42, 7 (July), 1043–1064.
- HOO, K. 2000. How much is enough? A risk-management approach to computer security. *Consortium for Research on Information Security Policy (CRISP) Working Paper*. Stanford University, Stanford, Calif., June.
- JAJODIA, S. AND MILLEN, J. 1993. Editors' preface. *J. Comput. Sec.* 2, 2/3, 85.
- JONES, A. 1997. Penetration testing and system audit. *Comput. Sec.* 16, 595–602.
- KPMG. 2000. Information Security Survey 2000. <http://www.kpmg.co.uk/services/audit/pubs/ISS> (Apr.), 1–4
- LARSEN, A. 1999. Global security survey: Virus attack. *InformationWeek.Com*. <http://www.informationweek.com/743/security.htm>.
- LITTLEWOOD, B., BROCLEHURST, S., FENTON, N., MELLOR, P., PAGE, S., WRIGHT, D., DOBSON, J., McDERMID, J., AND GOLLMAN, D. 1993. Towards operational measures of security. *J. Comput. Sec.* 2, 2, 211–229.
- LOCH, K. D., CARR, H. H., AND WARKENTIN, M. E. 1992. Threats to information systems: Today's reality, yesterday's understanding. *MIS Quart.* 17, 2, 173–186.
- LUOTONEN, O. 1993. Risk management and insurances. *Painatuskeskus Oy*. Helsinki, Finland.
- McKNIGHT, L., SOLOMON, R., REAGLE, J., CARVER, D., JOHNSON, C., GEROVAC, B., AND GINGOLD, D. 1997. Information security of internet commerce. In *Internet Economics*, L. McKnight and J. Bailey, Eds., MIT Press, Cambridge, Mass., pp. 435–452.
- MEADOWS, C. 2001. A cost-based framework for analysis of denial of service in networks *J. Comput. Sec.* 9, 1/2, 143–164.
- MILLEN, J. 1992. A resource allocation model for denial of service. In *Proceedings of the 1992 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, Los Alamitos, Calif., pp. 137–147.
- MURALIDHAR, K., BATRA, D., AND KIRS, P. 1995. Accessibility, security, and accuracy in statistical databases: The case for the multiplicative fixed data perturbation approach. *Management Sci.* 41, 9 (Sept.), 1549–1564.
- NIST (NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY). 1995. *An Introduction to Computer Security: The NIST Handbook*. (Special Publication 800-12).
- OSBORN, S., SANDHU, R., AND MUNAWER, Q. 2000. Configuring role-based access control to enforce mandatory and discretionary access control policies. *ACM Trans. Inf. Syst. Sec.* 3, 2 (May), 85–106.
- PEYRAVIAN, M., ROGINSKY, A., AND ZUNIC, N. 1999. Hash-based encryption. *Comput. Sec.* 18, 4, 345–350.
- PFLEEGER, C. 1997. *Security in Computing* (2nd ed.), Prentice-Hall, Englewood Cliffs, N.J.
- POWER, R. 2001. 2001 CSI/FBI computer crime and security survey. *Comput. Sec. J.* 17, 2 (Spring), 29–51.
- SANDHU, R. S., BHAMIDIPATI, V., AND MUNAWER, Q. 1999. The ARBAC97 model for role-based administration of roles. *ACM Trans. Inf. Syst. Sec.* 1, 2 (Feb.), 105–135.
- SANDHU, R. S., COYNE, E. J., FEINSTEIN, H. L., AND YOUMAN, C. E. 1996. Role-based access control models. *IEEE Comput.* 29, 2 (Feb.), 38–47.
- SCHNEIER, B. 1996. *Applied Cryptography* (2nd ed.), Wiley, New York.
- SIMMONS, G. 1994. Cryptanalysis and protocol failures. *Commun. ACM*, 37, 11 (Nov.), 56–64.
- STRAUB, D. W. 1990. Effective IS security: An empirical study. *Inf. Syst. Res.* 1, 3, 255–276.

- STRAUB, D. W. AND WELKE, R. J. 1998. Coping with systems risk: Security planning models for management decision making. *MIS Quart.* 23, 4, 441–469.
- VARIAN, H. R. 1997. How to build an economic model in your spare time. Part of a collection titled *Passion and Craft: Economists at Work*, ed. Michael Szenberg, University of Michigan Press, available at <http://www.sims.berkeley.edu/~hal/Papers/how.pdf>.
- VIGNA, G. AND KEMMEERER, R. A. 1999. NetSTAT: a network-based intrusion detection system. *J. Comput. Sec.* 7, 1, 37–71.
- WISEMAN, S. 1986. A secure capability computer system. In *Proceedings of the IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, Los Alamitos, Calif, pp. 86–94.

Received August 2001; revised May 2002; accepted June 2002