Information Security and Risk Management

by

Lawrence D. Bodin
Professor Emeritus of Decision and Information Technology
Robert H. Smith School of Business
University of Maryland
College Park, MD 20742

Lawrence A. Gordon
Ernst & Young Alumni Professor of Managerial Accounting and Information Assurance
Robert H. Smith School of Business
University of Maryland
College Park, MD 20742

Martin P. Loeb
Professor of Accounting and Information Assurance
Deloitte & Touche Faculty Fellow
Robert H. Smith School of Business
University of Maryland
College Park, MD 20742

Working Paper
Final Version Published in April 2008, *Communications of the ACM*, Vol. 51, no. 4, pp. 64-68

1. INTRODUCTION

Gordon and Loeb (2001, 2002, 2006b) provided an economic framework for evaluating information security activities. A key concept related to such a framework is the notion of risk management. Even though organizations try to avoid costly information security breaches, organizations cannot make their information 100% secure all of the time. Thus, managing the risk associated with potential information security breaches is an integral part of resource allocation decisions associated with information security activities.[1] Of course, to make resource allocation decisions, one needs to be clear on what is meant by the term *risk*.

Risk has a variety of meaning in the context of information security. The objectives of this paper are to discuss three measures that capture different aspects of information security risk and to propose a methodology that allows decision-makers to combine these (or any) different risk measures into a single composite metric. We call our proposed new metric the *Perceived Composite Risk (PCR).*

We suggest using the Analytic Hierarchy Process (Saaty, 1980) to determine the weighting factors that are used to combine these risk measures into the PCR. We also provide an example to show how decision-makers can use the PCR to evaluate different proposals for enhancing an organization's information security system. This paper builds on the Analytic Hierarchy Process (AHP) analysis provided by Bodin, Gordon and Loeb (2005) for assisting a Chief Information Security Officer (CISO) in ranking proposals for enhancing an organization's information security system. [2]

[1] See Gordon, Loeb, and Sohail (2003) for a framework of cyber risk management incorporating the use of insurance.

[2] Other papers addressing issues related to information security that are relevant to this paper are Gordon and Loeb (2006a) and Gordon, Loeb and Lucyshyn (2003).

2. THE PERCEIVED COMPOSITE RISK (PCR) METRIC

Three Measures of Risk

Three measures capturing commonly considered facets of risk are the expected loss, the expected severe loss, and the standard deviation of the loss. We now discuss each of these facets of risk.

The *expected loss* is derived by taking the sum of the product of each loss with its respective probability.[3] The expected loss is conceptually equivalent to the popular Annual Loss Expectancy (ALE) measure (see, for example, Gordon and Loeb, 2006b, pp. 75-79). Based on this metric, the larger the expected loss the larger the risk associated with an information security breach.

The *expected severe loss* focuses only on those breaches that put the survivability of the organization at risk. In order to calculate the expected severe loss, one has to first specify the magnitude of loss that, were it to occur, would threaten the very survivability of the organization. The expected severe loss is derived by taking the sum of the product of each severe loss (i.e., each loss that is greater or equal to the specified threshold) with its respective probability. Based on this metric, the larger the expected severe loss the larger the risk associated with an information security breach.

The *standard deviation of loss* (which is the square root of the variance of loss) represents the dispersion around the expected loss. The standard deviation of loss is computed by taking the square root of the product of squares of the deviation of each loss from the expected loss multiplied by the probability of that loss. Based on this metric, the larger the standard deviation, the larger the risk associated with a security breach. We used the standard deviation of

[3] For expositional ease, we assume that loss is a discrete random variable.

loss rather than the variance of loss because the standard deviation of loss is measured in the same units (dollars, for example) as the expected loss and expected severe loss.

The preceding three risk metrics can be illustrated with an example. Let X be a random variable representing the loss (in millions of dollars) attributable to a breach. Suppose for a proposal (called Proposal 1) for enhancing information security activities, X has the following discrete uniform distribution:

P[X=x] = .1 for x = 0, 1, 2, … , 9.

Therefore, the expected loss from a breach, E[X], under Proposal 1 is given by:

$$E[X] = \sum_{x=0}^{9} x \cdot P[X = x] = 0 \cdot [.1] + 1 \cdot [.1] + ... + 9 \cdot [.1] = 4.5$$

In order to calculate the expected severe loss, the decision-maker must first specify a threshold level. Suppose the threshold level, denoted by T, is judged to be 8, i.e., any breach whose cost is $8 million or greater is believed to put the survivability of the organization at risk. The expected severe loss under Proposal 1, denoted by E[severe loss] , is given by:

$$\text{E[severe loss]} = \sum_{x=8}^{9} x \cdot P[X = x] = 8 \cdot [.1] + 9 \cdot [.1] = 1.7$$

The standard deviation of loss, denoted by σ, under the loss function defined for Proposal 1 is given by:

$$\sigma = \sqrt{\sum_{x=0}^{9} (x - E[X])^2 \cdot P[X = x]} = \sqrt{8.25} \approx 2.872$$

We now present the PCR metric.

Computing the Expected Perceived Composite Risk (PCR)

For a given set of information security activities, the PCR is a linear combination of the expected loss, the expected severe loss, and the standard deviation of loss that can be attributable to a breach. Specifically,

$$PCR \equiv E[X] + [B/A] \cdot E[\text{severe loss}] + [C/A] \cdot \sigma$$

where the weights A, B, and C are determined from the AHP. The weights, A, B, and C are positive, sum to one, and reflect the relative importance of the performance metrics to the decision maker. An overview of the AHP (in an information security investment context) is given in Bodin, Gordon, and Loeb (2006).

Before turning to the question of how these weights are derived using AHP, we summarize the properties of the PCR:

- PCR is equal to the expected loss plus two penalty terms.
- The penalty term, $[B/A] \cdot E[\text{severe loss}]$, measures an additional perceived loss due to a severe loss occurring.
- The penalty term, $[C/A] \cdot \sigma$, measures an additional perceived loss due to variability in predicting the loss.

We now return to the interpretation and determination of the weights A, B, and C using AHP. The weights, A, B, and C, measure the emphasis that the CISO wants to place on the three risk measures (i.e., the expected loss, the expected severe loss, and the standard deviation). The weights on the three terms are 1, B/A and C/A. Without loss of generality, one can normalize the weights on the terms in the PCR so that the weight on the expected loss, E[X], is equal to one. In that way, if the user wants the PCR to equal the expected loss, the user would set B=0 and C=0 in the above equation defining PCR.

To illustrate the AHP method for determining the values of the weights, we consider an example (used in the next section). Table 1 presents the following pairwise comparison matrix of the three criteria - expected loss, expected severe loss, and the standard deviation of the loss. The pairwise comparison matrix is made up of columns 2-4 and rows 2-4 in Table 1. The final column in Table 1 gives the weights as determined by the eigenvector associated with the maximum eigenvalue for the pairwise comparison matrix given in columns and rows 2-4 in Table 1 (for further details, see Bodin, Gordon, and Loeb 2006).

	Expected Loss E[X]	**Expected Severe Loss**	**Standard Deviation of Loss σ**	**Weights**
Expected Loss E[X]	1	1	2	**.4**
Expected Severe Loss	1	1	2	**.4**
Standard Deviation of Loss σ	1/2	1/2	1	**.2**

Table 1: Pairwise Comparison Matrix and Weights for the Example

In establishing this pairwise comparison matrix, the assumption in this example is that the Expected Loss (E[X]) and Expected Severe are equally important criteria and both of these criteria are slightly more preferred to the Standard Deviation of Loss (σ) criterion. The pairwise comparisons that represent this judgment are realized by setting a_{12}=1, a_{21}=1, a_{13}=2, a_{23}=2, a_{31}=1/2, and a_{32}=1/2. Further, the diagonal elements, a_{11}, a_{22}, and a_{33} are set = 1, since a criterion is equally important to itself.

For a given decision-maker (e.g., the CISO), for which AHP reveals these weights (i.e., A=.4, B=.4 and C=.2), the value of the PCR for Proposal 1 is as follows:

$$\text{PCR}(\text{Proposal 1}) = 4.5 + [.4/.4]\cdot[1.7] + [.2/.4]\cdot[2.872] = 4.5 + 1.7 + 1.436 = 7.636$$

We now provide an example of how the PCR can be used in making information security investment decisions, where three different proposals for security activities exist.

3. EVALUATING FOUR PROPOSALS USING THE PCR

In order to demonstrate the use of the PCR, assume that the CISO is faced with selecting among four equal cost proposals for enhancing the organization's information security. Suppose that the CISO and his staff have estimated the loss probabilities associated with the three proposed sets of information security activities. The estimated loss probabilities associated with each proposal have been broken down into ten discrete amounts as displayed in Table 2.

	Losses from Information Security Breach in millions of dollars										
	0	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**Other values**
Probability of Loss-Proposal 1	.1	.1	.1	.1	.1	.1	.1	.1	.1	.1	0
Probability of Loss-Proposal 2	0	0	.2	0	0	.5	0	.1	.2	0	0
Probability of Loss-Proposal 3	.3	.2	0	0	0	0	.05	.05	.1	.3	0
Probability of Loss-Proposal 4	.0	.0	0	0	0	0	0	.45	.45	.1	0

Table 2: Probability of Losses under Three Information Security Project Proposals

We also continue to suppose that the threshold level, T, of a severe loss is $8 million. Table 3 shows the values of the three individual risk measures for each of the 3 proposals. Further, Table 3 gives the value of the PCR for each proposal, assuming that A = .4, B= .4 and C=.2.

	Expected Loss E[X]	**Expected Severe Loss**	**Standard Deviation of Loss σ**	**Perceived Composite Risk PCR**
Proposal 1	4.5	1.7	2.872	7.636
Proposal 2	5.2	1.6	1.990	7.795
Proposal 3	4.35	3.5	4.028	9.864
Proposal 4	7.65	4.5	0.654	12.477

Table 3: Risk Measures for the Three Proposals (where T=8, A=.4, B=.4 and C=.2)

Some of the problems with using the popular metric of`expected loss as a sole measure of risk can be easily seen by examining Tables 2 and 3. According to the expected loss metric, the

most preferred proposal is Proposal 3 to be followed in order by Proposal 1, Proposal 2 and Proposal 4. Note that although Proposal 3 minimizes the expected loss it also generates the second highest probability of threatening the survivability of the organization (Pr [X≥8] = .4) and generates the highest standard deviation of loss.

From Table 3, we also see that based on the expected severe loss criterion, the most preferred proposal is Proposal 2 to be followed in order by Proposal 1, Proposal 3 and Proposal 4. Further, based on the standard deviation criterion, Proposal 4 is most preferred proposal followed in order by Proposal 2, Proposal 1, and Proposal 3. Thus, a decision-maker interested in minimizing the risk of a breach, could rationally select Proposal 2, Proposal 3, or Proposal 4, depending on the **single** risk metric being considered.

The PCR combines the three risk measures through a procedure that carefully determines the decision maker's relative weighting of the risk criteria. The weights are decision maker dependent, so that the rankings based on the PCR may vary from person to person. With the values of A, B, and C given by .4, .4 and .2, respectively, Proposal 1 is preferred to Proposal 2, which in turn is preferred to Proposal 3, which is preferred to Proposal 4. It is interesting to note that Proposal 1 had the smallest value of the PCR even though Proposal 1 did not dominate any individual metric. However, if the decision maker's weights were A=.1, B=.2, and C=.7, then based on the PCR, Proposal 4 is preferred to Proposal 2, which is preferred to Proposal 1, which in turn is preferred to Proposal 3.[4]

Quite simply put, the common approach of using expected loss of a breach as the ranking criterion gives the CISO a narrow analysis of the alternatives and may lead to misleading results. Examining these other risk measures helps the CISO determine the best proposal to select and

[4] In this case, PCR(Proposal 4) = 21.227, PCR(Proposal 2) = 22.330, PCR(Proposal 1) =28.006, and PCR(Proposal 3) = 39.548.

implement. Although we formed the PCR as a linear combination of expected loss, expected severe loss, and standard deviation of loss, the method of forming a single PCR type of metric from a set of criteria is a general methodology. The decision-maker can use any set of criteria to form a PCR type of metric and use the AHP to determine the weighting factors. In that way, no matter what aspects of risk a decision-maker wishes to consider, a PCR type of metric can be a powerful decision-making tool.

4. FINAL REMARKS

Anyone responsible for information security must understand how to manage risk. Yet, the initial step of defining risk is far from easy. Popular measures of risk, such as expected loss from a breach or the standard deviation of a loss from a breach, only capture narrow facets of risk. In this paper, we introduced a new metric, called the Perceived Composite Risk (PCR), to evaluate the investment proposals for enhanced information security and suggested using the AHP to determine the weights in the PCR. The PCR gives the user some powerful new tools in analyzing proposals for enhancing an organization's information security system. Further, this analysis complements the analysis of Bodin, Gordon and Loeb [2005], which details how to effectively spend an information security budget, taking into account both non-financial and financial aspects of proposed information security projects.

REFERENCES

Bodin, L., L. A. Gordon, and M. P. Loeb. 2005. Evaluating Information Security Investments using the Analytic Hierarchy. *Communications of the ACM*, 48, No. 2 (February): 78-83.

Gordon, L. A. and M. P. Loeb. 2001. A Framework for Using Information Security as a Response to Competitor Analysis Systems. *Communications of the ACM* 44, No, 9 (September): 70-75.

Gordon, L. A., and M. P. Loeb. 2002. The Economics of Investment in Information Security. *ACM Transactions on Information and System Security* 5, No. 4 (November): 438-457.

Gordon, L. A. and M. P. Loeb. 2006a. Budgeting Process for Information Security Expenditures: Empirical Evidence. *Communications of the ACM* 49, No. 1 (January): 121-125.

Gordon, L. A. and M. P. Loeb. 2006b. *Managing Cybersecurity Resources: A Cost-Benefit Analysis*. New York: McGraw-Hill.

Gordon, L.A., M.P. Loeb, and W. Lucyshyn. 2003. Sharing Information on Computer Systems: An Economic Analysis. *Journal of Accounting and Public Policy* 22, No. 6: 461-485.

Gordon, L. A., M.P. Loeb and T. Sohail. 2003. A Framework for Using Insurance for Cyber Risk Management. *Communications of the ACM* 46, No 3 (March): 81-85.

Saaty, T. L. 1980. *The Analytic Hierarchy Process*. New York: McGraw-Hill.