

Working Paper

**How Much Is Enough?
A Risk-Management
Approach to Computer Security**

Kevin J. Soo Hoo

Consortium for Research on Information Security and Policy (CRISP)

CRISP is a research collaboration of the Center for International Security and Cooperation, the Institute for International Studies, and the School of Engineering, Stanford University.

June 2000

The opinions expressed here are those of the author and do not necessarily represent the positions of the Center, its supporters, or Stanford University.

© Copyright by Kevin J. Soo Hoo 2000
All Rights Reserved

Preface

The research for this working paper was sponsored in part by the Consortium for Research on Information Security and Policy (CRISP).

CRISP was created at Stanford University to develop a better analytical and policy understanding of national and international security problems relating to information technology. CRISP studies explore the technological, legal, organizational, and policy dimensions of these problems. The consortium includes university researchers from the Center for International Security and Cooperation and from two departments in the School of Engineering: the Department of Computer Science and the Department of Management Science and Engineering. CRISP works with companies involved in various areas of information technology, with network users and providers, and with parts of the federal government.

The specific projects undertaken by CRISP draw on the interests and knowledge of this community. The three main areas of work are a university/industry/government forum, technology and policy research, and international participation. CRISP's main function is to provide a forum to continue and expand the dialogue among the main stakeholders in U.S. national information infrastructures (i.e., the infrastructure owners, the network technology industry, the major users, the federal government, and the research community). CRISP members will continue to assist in the process of developing common views among these interested organizations through analysis of the surrounding issues.

In the technology and policy area CRISP defines and conducts research projects on subjects that are important to understanding the vulnerability of information infrastructures, the barriers to solutions, and possible remedies. These projects investigate and analyze technical constraints on infrastructure protection and possible technological developments, international policy considerations in protecting infrastructure, and the effect of existing and proposed laws and regulations on the goal of securing infrastructure.

Information infrastructure security is a manifestly international problem since usage, and hence dependence, are becoming global. Cyber attacks can move easily across borders, and adequate remedies will require a high degree of interstate cooperation. CRISP will, through conferences and other forms of exchange, undertake to build an international constituency to address the problems of securing information infrastructures on a global basis.

As a product of the technology and policy research area, this paper examines the resource allocation dilemma facing every organization that uses information technology: How much security is enough? The answer is found by investigating the applicability and utility of risk-management tools and techniques to computer-security risks. The paper begins with a short history of computer security risk management, highlighting the challenges and successes that marked each generation of risk-management tools. Next, it gives a brief discussion of the state of publicly available computer security data with recommendations for how that state might be improved. Finally, the paper offers a demonstration of a decision-analysis-based approach for managing computer security risks that directly addresses many of the issues that stymied previous computer security risk-management efforts.

Because much of the information infrastructure is privately owned and operated, efforts to improve general infrastructure security must be mindful of the resource allocation predicament confronting individual firms and organizations. By understanding the

economic pressures and incentives under which these actors formulate their individual security policies, public policymakers will be better able formulate national initiatives that supplement and enhance existing private security efforts.

This working paper is a published version of Kevin J. Soo Hoo's doctoral dissertation for the Department of Management Science and Engineering at Stanford University. For more information about CRISP and its activities, see its web page under research at <http://cisac.stanford.edu>.

Michael M. May, Senior Fellow
Institute for International Studies

Seymour E. Goodman, Director
Consortium for Research on Information Security and Policy

Abstract

How much security is enough? No one today can satisfactorily answer this question for computer-related risks. The first generation of computer security risk modelers struggled with issues arising out of their binary view of security, ensnaring them in an endless web of assessment, disagreement, and gridlock. Even as professional risk managers wrest responsibility away from the first-generation technologists, they are still unable to answer the question with sufficient quantitative rigor. Their efforts are handicapped by a reliance on non-quantitative methodologies originally developed to address the deployment and organizational acceptance issues that plagued first-generation tools.

In this dissertation, I argue that these second-generation approaches are only temporary solutions to the computer security risk-management problem and will eventually yield to decision-focused, quantitative, analytic techniques. Using quantitative decision analysis, I propose a candidate modeling approach that explicitly incorporates uncertainty and flexibly allows for varying degrees of modeling detail to address many of the failings of previous modeling paradigms. Because quantitative modeling requires data, I also present a compilation and critique of publicly available computer security data. I highlight the importance of data collection, sharing, and standardization with discussions of measurement, relevance, terminology, competition, and liability. I conclude with a case study example, demonstrating how uncertain data and expert judgments are used in the proposed modeling framework to give meaningful guidance to risk managers and ultimately to answer the question: How much is enough?

Contents

Preface	ii
Abstract	iv
Chapter 1 Introduction and Background	1
1.1 Information Security	2
1.2 Risk Assessment and Risk Management.....	3
1.3 Common Framework.....	4
1.4 What Went Wrong.....	7
1.5 Second-Generation Approaches	9
1.5.1 <i>Integrated Business Risk Management Framework</i>	<i>9</i>
1.5.2 <i>Valuation-Driven Methodologies</i>	<i>10</i>
1.5.3 <i>Scenario Analysis Approaches.....</i>	<i>11</i>
1.5.4 <i>Best Practices</i>	<i>11</i>
1.6 Underlying Forces for Change	12
Chapter 2 Risk Modeling and Analysis.....	15
2.1 Risk Modeling As a Decision-Driven Activity.....	15
2.2 Decision Modeling.....	16
2.3 Computer Security Risk Model Description	19
2.4 Analysis Techniques.....	23
2.5 Summary.....	27
Chapter 3 Data.....	29
3.1 Difficulty with Data	29
3.2 Diffidence about Data Sharing.....	31
3.3 Relevance of Data	32
3.4 Finding Data.....	33
3.4.1 <i>Terminology</i>	<i>35</i>
3.4.2 <i>Annual Frequencies of Security Incidents.....</i>	<i>38</i>
3.4.3 <i>Consequences of Security Incidents</i>	<i>40</i>
3.4.4 <i>Safeguards</i>	<i>43</i>
3.5 Summary.....	46
Chapter 4 Example Model and Analysis	47
4.1 First Iteration.....	47
4.1.1 <i>Parameters</i>	<i>47</i>
4.1.2 <i>Input Variables</i>	<i>49</i>
4.1.3 <i>Initial Results.....</i>	<i>53</i>
4.1.4 <i>Analysis</i>	<i>54</i>
4.2 Subsequent Iterations	60
4.2.1 <i>Consequences of Bad Events.....</i>	<i>60</i>
4.2.2 <i>Frequency of Bad Events</i>	<i>62</i>
4.2.3 <i>Costs of Safeguards</i>	<i>63</i>
4.3 Model Adaptability.....	64

Chapter 5	Conclusions and Implications	67
Appendix A	Analytica Code for Example Model	70
Bibliography	81

List of Tables

Table 1. Percentage of Respondents Reporting One or More Security Incidents	39
Table 2. CSI/FBI Reported Financial Losses Due to Security Incidents	42
Table 3. Safeguards Selector	51
Table 4. Safeguard Reductions in Bad Event Frequencies	52
Table 5. Safeguard Reductions in Bad Event Consequences.....	52
Table 6. Annual Frequency of Bad Events without New Safeguards	53
Table 7. Consequences of Bad Events without New Safeguards.....	53

List of Figures

Figure 1. Common Framework Process Diagram	5
Figure 2. Common Framework Event Tree Example	8
Figure 3. Building Blocks of Influence Diagrams	17
Figure 4. Example of a Decision Diagram	18
Figure 5. Computer Security Risk Management Decision Diagram with Variable Identifiers	20
Figure 6. Stochastic Dominance in Cumulative Probability Distributions	25
Figure 7. Sets of Security Breach Attempts	30
Figure 8. Vulnerability Analysis & Assessment Program Results, 1996	31
Figure 9. Survey Respondents' Information Security Budgets	34
Figure 10. ASIS and CSI/FBI Survey Population Characteristics	34
Figure 11. CSI/FBI Results on the Detection and Number of Incidents	38
Figure 12. ASIS Report on Intellectual Property Theft	39
Figure 13. Information Week Report on Security Incident Consequences	43
Figure 14. CSI/FBI Report on Technologies in Use	44
Figure 15. Information Security Report on Technologies in Use	45
Figure 16. Information Week Report on Technologies in Use	45
Figure 17. Computer Security Risk Management Decision Diagram	48
Figure 18. Annual Loss Module Diagram	51
Figure 19. Expected Net Benefit for Three Policies	54
Figure 20. Net Benefit Cumulative Distribution	54
Figure 21. Safeguard Savings to Cost Ratios	55
Figure 22. Tornado Diagram.....	56
Figure 23. Cross-over Points for Either Initial Consequences or Initial Frequencies.....	57
Figure 24. Cross-over Points for Reduction in Frequency	58
Figure 25. Cross-over Points for Costs of Safeguards.....	58
Figure 26. Expected Value of Perfect Information	59
Figure 27. Detailed Diagram for Consequences of Bad Events	61
Figure 28. Detailed Diagram for Frequency of Bad Events.....	63
Figure 29. Detailed Diagram for Costs of Safeguards	64
Figure 30. Detailed Diagram for Liability and Embarrassment Consequences.....	65

Chapter 1 Introduction & Background

The revolutionary idea that defines the boundary between modern times and the past is the mastery of risk: the notion that the future is more than a whim of the gods and that men and women are not passive before nature. Until human beings discovered a way across that boundary, the future was the mirror of the past or the murky domain of oracles and soothsayers who held a monopoly over knowledge of anticipated events.

— Peter Bernstein, *Against the Gods: The Remarkable Story of Risk*¹

Peter Bernstein describes a fascinating rite of passage that Western civilization completed to enter the modern era. In many ways, this passage has not yet come to computer security. Since the dawn of modern computing, computer security has been left in the hands of “computer security experts,” chiefly technologists whose technical understanding qualified them to shoulder the responsibility of keeping computers and their valuable information safe. The rapid growth of society’s dependence upon information systems, the Internet being one of the most prominent examples, has precipitated a growing apprehension about the security and reliability of this fragile infrastructure. Recognizing that human behavior plays a significant role in computer security, often superseding the technological aspects, many organizations are shifting computer and information security responsibility away from computer security technologists and into the hands of professional risk managers. This change in the perception and practice of computer security is reminiscent of the first steps taken by Western civilization to end its dependence upon the ancient seers.

Although this transition might appear inevitable, practitioners of the “computer security folk art” are loath to concede without a fight. A fractious debate engulfs the computer security community, with rival factions and beliefs vying to shape the current and future practice of computer security risk management. This debate centers on the degree of quantitative rigor that risk assessment can reasonably achieve.

The variety of answers to the question “How much security is enough?” is an apt encapsulation of the differences separating the competing factions. From the best-practices approach to the quantification attempts of first-generation methodologies, none of the past or present approaches can satisfactorily support its answer to the question. The lack of computer security metrics and statistics precludes the grounding of their results in quantifiable reality. With the growing importance of the information infrastructure to the economy and to society as a whole, the security of that infrastructure has emerged as a leading national security and law enforcement issue. The seriousness of its implications means that the computer security risk management debate can no longer be confined to the elite community of computer security experts.

In this chapter, I set a context for the debate by tracing the evolution of computer security risk modeling. After a brief introduction to information security and risk assessment/management, I describe and critique the first generation of annual loss-expectancy-based risk models, typified by the common framework developed in the Computer Security Risk Management Model Builders Workshops. Next, I present four different approaches that constitute a second generation of risk models. I argue that the

¹ Peter L. Bernstein, *Against the Gods: The Remarkable Story of Risk* (New York: John Wiley & Sons, Inc., 1996), p. 1.

weaknesses in these second-generation models make them inadequate for the needs of the insurance industry, the calculus of legal liability, and the efficiency of competitive markets. As a result, a new, more quantitative approach will be needed to meet the future demands of risk modeling.

Against this backdrop, I present in Chapter 2 a candidate for such a new quantitative approach. Following the National Research Council's advice on risk management, the concept of a decision-driven analytic process is adapted to computer security risk modeling. The new framework incorporates the tools and analysis techniques of decision analysis and addresses many of the shortcomings of the previous risk modeling efforts.

In Chapter 3, I give an exposition and critique of publicly available data relevant to risk management. I discuss the importance of information sharing and how the lack of it continues to hamper risk assessment efforts. I also examine other issues relevant to data, including the need for standardization, the relevance of past data to future events, appropriate accounting of incident losses, and safeguard costs.

In Chapter 4, I marry some of the data presented in Chapter 3 to the approach outlined in Chapter 2 to demonstrate how the proposed framework would work. One full pass through the analysis cycle is presented and explained with suggestions for how the model might be further extended if warranted. In Chapter 5, I reflect on the implications of this research and suggest directions for future work.

1.1 Information Security

Information security is hardly a new concept. The need to protect valuable information is as old as mankind. Whether that information was the location of a rich hunting ground, a technology for producing weapons, or a knowledge of the divine, it had value and therefore needed protection. Today, information security is often conceptualized as being the protection or preservation of four key aspects of information: availability, integrity, authenticity, and confidentiality.

Availability:	Accessibility of information for a purpose.
Integrity:	Completeness, wholeness, and readability of information, and the quality of being unchanged from a baseline state.
Authenticity:	Validity, conformance, and genuineness of information.
Confidentiality:	Limited observation and disclosure of knowledge to only authorized individuals. ²

While the idea of information security is certainly not new, the practice of information security has been and continues to be an evolving endeavor wherein technological advances both help and hinder its progress. The advent of the information age, as heralded by the rapid and extensive diffusion of digital computing and networking technologies, is one such advance that has fundamentally changed the practice of information security by adding a new and dynamic dimension of computer security. The resultant growth in the volume of valuable information and the avenues by which it may be compromised has dramatically escalated the challenge of information security. Using computer systems and networks and capitalizing on weaknesses in equipment and human operators, malefactors

² Donn B. Parker, *Fighting Computer Crime: A New Framework for Protecting Information* (New York: John Wiley & Sons, Inc., 1998), p. 240. Parker's definitions were chosen for their brevity and because he is a stickler for language. Substantively similar definitions can be found elsewhere in the literature, such as Deborah Russell and G. T. Gangemi, Sr., *Computer Security Basics* (New York: Thunder Mountain Press, 1994) or John D. Howard, *An Analysis of Security Incidents on the Internet 1989-1995*, Ph.D. thesis, Department of Engineering and Public Policy, Carnegie Mellon University, April 7, 1997.

are able to strike at information assets with a whole host of attacks that twenty years ago were unimaginable, much less a daily reality.

The information technologies are a powerful set of enabling technologies. They confer upon their users an unprecedented capability for managing, processing, and communicating information. Securing these technologies and the information that they steward is a difficult and often expensive venture. In addition to the direct costs of planning, designing, and implementing safeguards, computer security also requires the participation of everyone in the organization and typically limits their freedom to use the technology to its fullest extent. Herein lies a fundamental tension between security and usability: security requires that information and access to it be tightly controlled whereas the key advantage of the information technologies is their ability to enable the free flow of information. In competitive industries, the outcome of this balancing act has been, predictably, struck in favor of usability over security.

The proponents of better security face an even more difficult task because good statistics on computer-related crime are rare. Advocates must convince their senior management to spend actual resources to address hypothetical losses. Because security initiatives are of uncertain and poorly quantified benefit, they could ultimately reduce computer usability and worker productivity while not providing any tangible benefits. Not surprisingly, these decisions favor security only when the security advocate commands significant respect from senior management. This dependence harkens back to an age when revered oracles dispensed wisdom about the future. If the hallmark of the modern age is the mastery of risk, then modern times demand that computer security risks, like many others, be managed with an open and rational process that depends more on quantified costs and benefits than on the pronouncements of a guru.

1.2 Risk Assessment and Risk Management

A formal risk framework can be a useful tool for decomposing the problem of risk management. In such a framework, risks are assessed by evaluating preferences, estimating consequences of undesirable events, predicting the likelihood of such events, and weighing the merits of different courses of action. In this context, risk is formally defined as a set of ordered pairs of outcomes (O) and their associated likelihoods (L) of occurrence.

$$\text{Risk} \equiv \{(L_1, O_1), \dots, (L_i, O_i), \dots, (L_n, O_n)\}^3 \quad (1)$$

Risk assessment is the process of identifying, characterizing, and understanding risk; that is, studying, analyzing, and describing the set of outcomes and likelihoods for a given endeavor. Modern risk assessment traces its roots to the nuclear power industry, where carefully constructed risk assessment methodologies were developed to analyze the operations of the very new and potentially dangerous nuclear power facilities. These methodologies centered around fault/event trees that were used to illustrate and to capture all possible plant failure modes in a graphical representation.

Risk management is a policy process wherein alternative strategies for dealing with risk are weighed and decisions about acceptable risks are made. The strategies consist of policy options that have varying effects on risk, including the reduction, removal, or reallocation of risk. In the end, an acceptable level of risk is determined and a strategy for achieving that level of risk is adopted. Cost-benefit calculations, assessments of risk tolerance, and quantification of preferences are often involved in this decision-making process.

³ Hiromitsu Kumamoto and Ernest J. Henley, *Probabilistic Risk Assessment and Management for Engineers and Scientists*, 2nd edition (New York: Institute of Electrical and Electronics Engineers, Inc., 1996), p. 2.

1.3 Common Framework

In 1979, the National Bureau of Standards published its Federal Information Processing Standard (FIPS) 65, *Guideline for Automatic Data Processing Risk Analysis*.⁴ The document set the risk assessment standard for large data-processing centers and also proposed a new metric for measuring computer-related risks: Annual Loss Expectancy (ALE).

$$ALE = \sum_{i=1}^n I(O_i)F_i \quad (2)$$

where :

$\{O_1, \dots, O_n\}$ = Set of Harmful Outcomes

$I(O_i)$ = Impact of Outcome i in dollars

F_i = Frequency of Outcome i

Although ALE was never itself enshrined as a standard, many treated it as such in subsequent work on risk-management model development.⁵ The metric's appeal rests in its combination of both risk components into a single number. Unfortunately, this blending of quantities has the disadvantage of being unable to distinguish between high-frequency, low-impact events and low-frequency, high-impact events. In many situations, the former may be tolerable while the latter may be catastrophic.

In the mid-1980s, the National Bureau of Standards (now a part of the National Institutes of Standards and Technology, or NIST) and National Computer Security Center (NCSC) seeded research in the area of computer security risk-management modeling.⁶ Using a series of workshops to focus the energies of risk experts specifically on the challenges of computer security risks, the organizations midwived the birth of a new and needed field of research. The methodologies and commercial software packages that sprang from this effort constitute the first generation of computer security risk-management models. Unfortunately, with the close of the decade the research activity all but ended. However, the work did generate a significant amount of attention during its lifetime and formed the basis of modern computer security risk assessment. Thus, when security experts, such as Donn Parker,⁷ decry the practice of risk assessment, they have a fairly specific methodology in mind, one that can be well understood by looking at the output of the NIST/NCSC workshops.

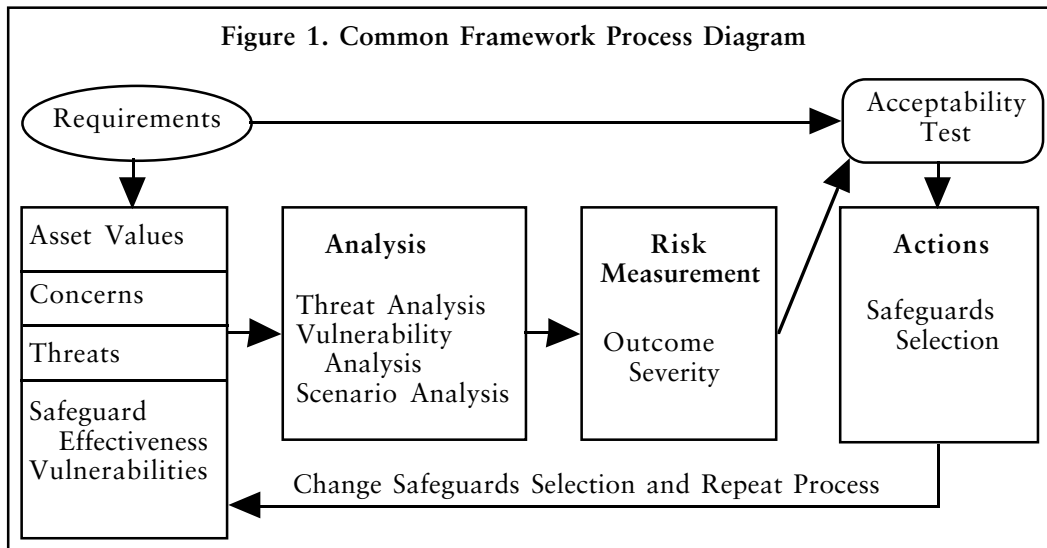
⁴ National Bureau of Standards, *Guideline for Automatic Data Processing Risk Analysis*, FIPS PUB 65 (Washington, DC: U.S. General Printing Office, 1979).

⁵ See the *Proceedings of the Computer Security Risk Management Model Builders Workshop* (Washington, DC: National Institutes of Standards and Technology, 1988) for several methodologies based on ALE. Among currently available commercial software packages, Bayesian Decision Support System from OPA Inc., Buddy System from Countermeasures, Inc., and CRAMM from International Security Technology implement ALE-based methodologies.

⁶ The workshops were called the Computer Security Risk Management Model Builders Workshops and were held over a two-year period from 1988 to 1989. The proceedings of the workshops are available from the National Institutes of Standards and Technology.

⁷ Donn B. Parker, *Fighting Computer Crime: A New Framework for Protecting Information* (New York: John Wiley & Sons, Inc., 1998), pp. 262–82.

By the time the workshops ended, a consensus framework for computer security risk management had emerged. Although never formalized as a NIST standard, the framework represents a good summary of the participants' collective thoughts.



The framework had seven basic elements:⁸

- Requirements: $R \equiv [R_1, R_2, \dots, R_i]$
e.g., expected loss < \$100K, expected loss < \$1M
- Assets: $A \equiv [A_1, A_2, \dots, A_k]$
e.g., hardware, software, data
- Security Concerns: $C \equiv [C_1, C_2, \dots, C_s]$
e.g., confidentiality, integrity, authenticity
- Threats: $T \equiv [T_1, T_2, \dots, T_m]$
e.g., human, natural
- Safeguards: $S \equiv [S_1, S_2, \dots, S_p]$
e.g., physical, system, communication, admin.
- Vulnerabilities: $V \equiv [V_1, V_2, \dots, V_q]$
e.g., physical, software, hardware, administrative
- Outcomes: $O \equiv [O_1, O_2, \dots, O_r]$
e.g., combinations of A, C, T, S, V

The framework also included three associated quantities:

- Asset Values: $A_{val} \equiv [A_{1val}, A_{2val}, \dots, A_{kval}]$
- Safeguard Effectiveness: $S_{eff} \equiv [S_{1eff}, S_{2eff}, \dots, S_{peff}]$
- Outcome Severity: $O_{sev} \equiv [O_{1sev}, O_{2sev}, \dots, O_{rsev}]$
e.g., ALE of the outcome, qualitative judgment

The framework called for an assessment of the above quantities in an iterative process as diagrammed in Figure 1. First, identification of security requirements, assets for consideration, security concerns, possible threats, vulnerabilities, and safeguards takes place. Next, a series of analyses ensues.

⁸ This description was adapted from Lance J. Hoffman and Brian T. Hung, "A Pictorial Representation and Validation of the Emerging Computer System Security Risk Management Framework," Computer Security Risk Management Model Builders Workshop, Ottawa, Canada, June 20–22, 1989, p. 6.

The threat analysis involves an examination of possible threats to each asset. The threats might include human actors, natural catastrophes, unintentional errors, etc. The vulnerability analysis looks at the weaknesses in security that might enable a successful attack against the assets. Although much agreement was reached while forging the common framework, the idea that vulnerabilities represented the absence of specific safeguards, rather than an independent variable in itself, remained controversial.⁹ The scenario analysis requires a detailed evaluation of assets, security concerns, threats, and vulnerabilities to generate all possible scenarios whereby security compromises could occur. These scenarios are then used in the risk-measurement phase to evaluate associated outcomes and to rate the magnitude of the risk. The acceptability test compares the risk measured for a given asset with the established requirements. Safeguard selection decisions are then made to close the gap between the required and measured risk levels. The entire process is then repeated under the new safeguard regime, resulting in a new risk measurement for each asset. These risk measurements along with assessments of safeguard costs are then used to generate cost-benefit analyses for each safeguard.

The common framework is quite generic in its specification and therefore broad in its potential application. The methodology can be adapted to either qualitative or quantitative risk assessment. The scenario analysis and subsequent risk-measurement activity are specifically well-suited to qualitative assessment. In a quantitative risk assessment, both tasks could be automated calculations, based on asset values, frequency of vulnerability exploitation, and probability of successful attack.

Several computer software applications implementing schemes similar to the common framework were commercialized during the late 1980s and early 1990s. A handful of applications, such as @Risk, BDSS, CRAMM, and the Buddy System, are still available today in varying stages of commercial release.¹⁰ Despite the best efforts of those involved, the common framework and other ALE-based approaches failed to gain widespread acceptance, and when the initial funding from NIST and NCSC dried up most of the risk modeling work ended with it.

Throughout the 1990s, sporadic research efforts in computer security risk models can be found, including proprietary projects done by large industrial companies for internal risk assessment,¹¹ computer surety work at Sandia National Laboratories,¹² and various efforts by individual security consultants and academics. Detailed descriptions of proprietary work are generally not publicly available. Most of the work on computer surety is available via electronic download, but many of the researchers have since moved on to other projects, unrelated to computer surety. The few publications that can be found in the recent open literature tend to focus on the deployment issues and do not, generally, break any new ground or resolve any other fundamental challenges of the common framework.¹³

⁹ The term “vulnerability” in information security today often refers to a specific weakness in hardware or software that enables security breaches. The correspondence between the safeguards that eliminate or reduce these weaknesses and the vulnerabilities themselves is not one-to-one. Thus, the lack of agreement in the late 1980s was perhaps prophetic of the current reality.

¹⁰ For more information on these and other automated risk-management tools from the early 1990s, see NIST, *Description of Automated Risk Management Packages That NIST/NCSC Risk Management Research Laboratory Has Examined*, Updated 1991 <<http://csrc.nist.gov/training/risktool.txt>> last accessed 27 October 1999.

¹¹ General Accounting Office, *Information Security Risk Assessment: Practices of Leading Organizations*, Exposure Draft <<http://www.gao.gov/monthly.list/aug99/aug991.htm>> last accessed 1 October 1999.

¹² Roxana M. Jansma, Sharon K. Fletcher, et al., *Risk-Based Assessment of the Surety of Information Systems* (Springfield, VA: National Technical Information Service, 1996).

¹³ For examples, see Fred Cohen, “Managing Network Security: Balancing Risk,” December 1998, <<http://all.net/journal/netsec/9812.html>> last accessed 14 November 1999; Charles Cresson Wood, “Using Information Security to Achieve Competitive Advantage,” *Proceedings of the 18th Annual Computer Security*

1.4 What Went Wrong

In retrospect, three fatal flaws doomed the common framework and its ALE-based brethren to failure. The deficiencies are as much a reflection of the inventors' biases as they are an illustration of the challenges that face any attempt to model computer security risks.

First, the methodology's scenario-generation mechanism created an assessment task of infeasible proportions. In any mathematical or computer modeling endeavor, a balance must be struck between model simplicity and faithful replication of the modeled system. If the model errs on the side of simplicity, then it may not be sufficiently accurate to be of any use. If, on the other hand, it errs on the side of faithful replication, then its implementation may be so overwhelming as to render it impracticable. This tension pervades every modeling effort. Unfortunately, the ALE-based methodologies tended to favor significantly greater detail than was efficiently feasible to describe.

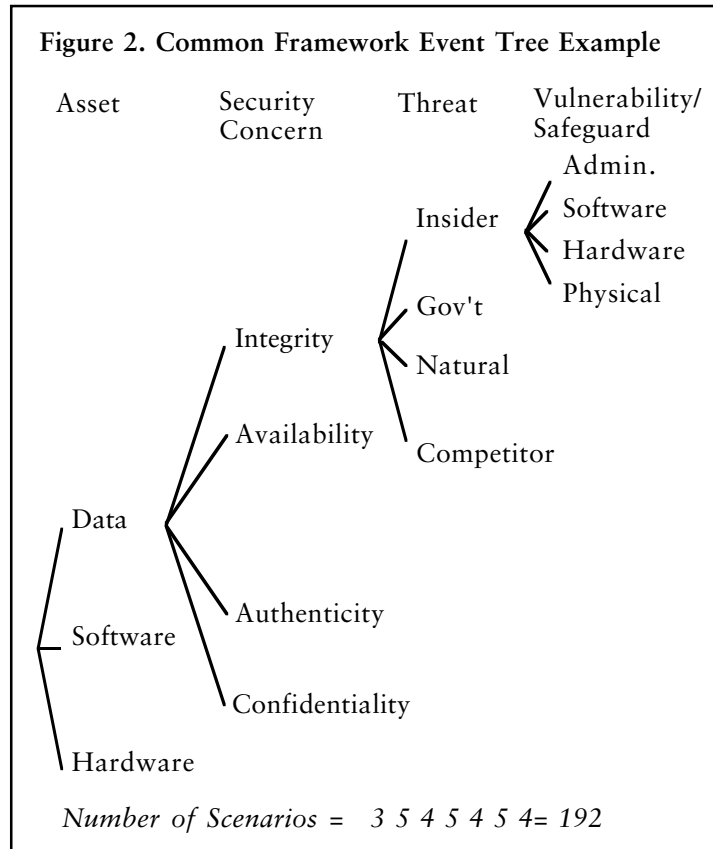
To illustrate the point, Figure 2 presents an event tree similar to one that might have been constructed during the scenario-analysis phase of the common framework. In this simple example, only three or four items are listed at each branching point, and each of those items is merely a general category summarizing many more specific items. Each path through the tree constitutes a scenario that includes an asset, a security concern, a threat, and a vulnerability/safeguard. Once the tree is fully specified, these scenarios must be analyzed and evaluated for their frequencies of occurrence and potential impacts.

When experts attempted to apply methodologies like the common framework, they quickly found themselves mired in an immense assessment task with thousands to hundreds of thousands of branches. This result came about because the number of scenarios grows almost exponentially with each new item considered, i.e., the number of scenarios $\approx k \times s \times m \times p$.¹⁴ This result is also a product of the technologists' binary view of security. In their eyes, systems are either secure, in which case they have no vulnerabilities, or are insecure, in which case they have vulnerabilities that require remedial action. In an effort to be comprehensive in their risk assessment, these early risk modelers tried to capture all threats, assets, vulnerabilities, and security concerns, from the most common to the most obscure. This conscientiousness and the ensuing exponential scenario growth quickly put the assessment, data collection, and analysis requirements well beyond the capabilities of both personnel and computing resources.¹⁵ This prohibitively expensive assessment process was one of the pivotal reasons that ALE-based risk models failed to achieve widespread acceptance and implementation.

Institute Conference, Miami, November 11–15, 1991; The Economist Intelligence Unit, *Managing Business Risks in the Information Age* (New York: The Economist Intelligence Unit, 1998); Amit Yoran and Lance C. Hoffman, "Role-based Risk Analysis," *Proceedings of the 20th National Information Systems Security Conference*, October 1997, p. 587.

¹⁴ The number of scenarios is actually less than or equal to the product because not all concerns, threats, or vulnerabilities are applicable to every asset.

¹⁵ Bear in mind that when this methodology was being touted, the fastest desktop computers ran on the Intel 80386 or the Motorola 68020 microprocessors, roughly four microprocessor generations ago.



The second flaw of the first-generation models is also rooted in the technologists' binary view of security. These models were completely deterministic, assuming that all quantities would be precisely known. Variables were assessed as single-point estimates rather than as probabilistically weighted or parameterized ranges of values. The inability to recognize and capture uncertainty handicapped risk modelers. Quite often, these estimates became sources of great controversies that distracted organizations and derailed security efforts. In the end, the methodology was blamed, and both organizations and computer security experts were reluctant to use it again.

The last significant challenge facing the common framework is its dependence on information that was and continues to be extremely sparse. A model is only as good as the information put into it. The common framework, like all ALE-based methodologies, assumes that frequency, valuation, and efficacy data are available. In the computer security arena today, such data remains largely unavailable. In fact, no highly effective automated tools have been developed to measure safeguard efficacy or to record computer security breaches consistently. As for valuing information and determining the consequences of a security breach, no methodologies have been standardized.

The lack of tools and standards is very likely a by-product of an institutional predisposition against the compilation of security statistics. Ignorance and fear perpetuate this attitude.¹⁶ Ignorance is a self-reinforcing problem since organizations are reluctant to act on security concerns unless a real problem has been proven. Such proof, however, will not be forthcoming without an initial investment in security for monitoring purposes. Hence, a chicken-and-egg problem ensues until a strong advocate is able to secure the

¹⁶ Ironically, fear is also the tool that is used most by computer security advocates to convince their clients that implementing better security practices now will be less painful than the consequences of security breaches later.

initial investment without quantitative proof or the consequences of insecurity are so significant as to be recognizable without formal monitoring activities.

Fear of lawsuits from shareholders or business partners due to inadequate security motivates some companies to avoid collecting security statistics altogether. If a company were to collect such statistics and suffered a breach of security, resulting in a lawsuit, then those statistics could be subpoenaed and used against the company to potentially demonstrate negligence. If an analysis of the statistics showed that the company was not only aware of security lapses but also did not take sufficient precautionary measures, the company could be found negligent and therefore liable for damages. To avert this chain of events, some companies simply refuse to track security incidents. Once again, this blissful ignorance is only a transient condition. If, as security advocates would have us believe, computer security breaches are or will soon be too significant to be ignored, then the fact that the company does not keep computer security records could in itself be grounds for a lawsuit alleging management negligence.

Excessive complexity, poor treatment of uncertainty, and data unavailability spawned implementation and deployment challenges that were beyond the capabilities of common framework proponents to address. Although the ALE-based methodologies provided risk assessors with a blueprint for how to proceed, no concrete examples were ever published to demonstrate their application. As a result, many organizations turned a blind eye to computer security risk modeling and management. This condition might have persisted to this day were it not for the advent of the Internet and the lucrative promise of e-commerce.

1.5 Second-Generation Approaches

The world is a much changed place since the late 1980s. Four desktop microprocessor generations have passed; the typical desktop computer system memory has increased by an order of magnitude or more; and the explosion of the Internet has fundamentally changed the ways in which people view and use computers. These changes have both encouraged and enabled a renewed interest in computer security risk management. Security is no longer viewed exclusively as an inconvenient, additional cost, but as a key for unlocking new business opportunities, such as electronic commerce ventures. Along with this revived interest have come new approaches to managing risk, with new champions and advocates. Although the technologists who led the first efforts still remain, their role is somewhat diminished, having become one of technical assistance and support. The new approaches tend to focus almost exclusively on solving the deployment and organizational acceptance issues that plagued the earlier ALE generation, thus leaving the challenges of complexity and uncertainty unaddressed. Most of the new and innovative work is hidden behind a veil of corporate secrecy, rendering a full exposition and analysis of the second-generation approaches unachievable; however, enough open-source material exists to present a basic description and critique of four general approaches from leading organizations, including the Integrated Business Risk-Management Framework, Valuation-Driven Methodologies, Scenario Analysis Approaches, and Best Practices.

1.5.1 *Integrated Business Risk-Management Framework*

The Integrated Business Risk Management concept stems from the idea that information-technology-related risks are like any other serious “business risk” and must therefore be managed as such. Business risks are broadly defined to include operational risks, financial risks, environmental risks, and others. This conceptual breakthrough provides a new context for computer security risks, larger than the immediate hardware, software, and information that typically made up a risk assessment. As a consequence, the implementation and assessment tasks could become even more daunting. However, casting

these risks in the same light as others forces simplifications in assessment and analysis. The modeling emphasis moves away from capturing the details of computer security interactions, focusing instead on bottom-line business impact, or value added.¹⁷

The Integrated Business Risk Management model is championed by management consulting and computer security strategy consulting firms. By pointing out the critical role of information technologies in business processes, they are able to demonstrate the necessity of security and to highlight its importance to current and future business opportunities. The methodology is distinctly non-technical and typically motivates a subsequent discussion about the specific safeguards that should be taken to improve security. This second phase of detailed risk assessment must be performed after the integrated business risk management exercise to formulate a plan of action for securing information assets. Some examples of businesses that practice integrated business risk management include Microsoft, Mitsui, Capital One Financial, Fidelity Management and Research, and BOC Gases Australia.¹⁸

1.5.2 Valuation-Driven Methodologies

The impracticality of ALE-based methodologies, with their massive assessment needs, forced risk managers to develop alternatives that would be less prone to controversy and more easily implemented, the dearth of data notwithstanding. Recalling that risk is made up of consequences and their respective likelihoods, or frequencies, of occurrence and that no sufficiently detailed statistics are available to predict those likelihoods, a reasonable simplification might be to ignore the likelihood half of the risk definition. Valuation-driven approaches make just such a simplification.

Unlike the high-level, integrated business risk-management methodologies described above, valuation-driven techniques are quite specific in their safeguards recommendations. These specifications, assigned based upon asset value alone, are intended both to ensure security and to standardize security practices throughout an organization.

The process of implementing a valuation-based methodology is fairly straightforward. Once security policies have been drawn up for each asset-value category, the assets must be inventoried and valued. The responsibility for determining an asset value must be prudently assigned to an individual who possesses a broad view and an understanding of the asset's role in the organization. After the valuation is made, the corresponding security specification is produced, and a plan is formulated to meet the specification. Most of the technical challenges that hamstrung the first-generation methodologies are avoided with this valuation-based approach, allowing implementers to focus on issues of managerial acceptance, institutional inertia, and other deployment issues.

Although attractive for their simplified approach and avoidance of controversy, these valuation-driven methods suffer significant theoretical flaws by virtue of that same simplification. Their exclusive focus on asset value and ignorance of safeguard costs, efficacy measures, and frequency of security breaches could result in either over-securing assets or under-securing them. Both possibilities are economically inefficient and could cause competitiveness to suffer. Without the capacity for performing cost-benefit analysis or the requirement of basic statistics collection, these methods provide no mechanism to motivate refinement of their security specifications. Although convenient in the short term, valuation-driven approaches are not viable long-term solutions. Nevertheless, several companies have begun implementing these programs as a first step toward standardizing

¹⁷ The term "value added" is often used in this context to describe the positive effect that a resource, process, or policy has on an organization.

¹⁸ For a more complete description of the Integrated Business Risk-Management model, including case study examples, see The Economist Intelligence Unit, *Managing Business Risks in the Information Age* (New York: The Economist Intelligence Unit, Ltd., 1998).

and improving their corporate computer security, including Sun Microsystems, International Business Machines, and JP Morgan.¹⁹

1.5.3 Scenario Analysis Approaches

Scenario-analysis approaches are probably more common than any others, especially in small-to-medium sized enterprises. As its name implies, scenario analysis involves the construction of different scenarios by which computer security is compromised. Scenario analysis is often employed to dramatically illustrate how vulnerable an organization is to information attacks. For example, some consultants will, with their client's permission, hack into the client's information systems, obtain sensitive data, and provide the client with a report detailing the data stolen, how quickly it was obtained, and other particulars of the exploit. This "red-teaming" exercise helps motivate the client to pursue better security and to provide further work for security consultants.

Scenario-analysis techniques are also used to encourage broader brainstorming about computer-related risks. Some companies have small information technology risk management teams whose experts fan out to company divisions, provide facilitation services and technical expertise, and help the divisions understand their risk exposure. In this setting, scenarios are used to demonstrate the variety and severity of risks faced. The scenarios deemed most likely and of greatest severity are then used as the basis for developing a risk-mitigation strategy.

The primary drawback of an exclusively scenario-analysis approach is its limited scope. Looking back at the event tree example from the common framework in Figure 2, scenarios essentially represent different paths through the tree. The danger of assessing only a few scenarios is the possibility that important paths may be missed, leaving serious risks unaddressed. By narrowing the focus in this way, the analysis is made tractable, but incomplete. In addition, this approach also does nothing to encourage a better, more comprehensive data collection activity. Like the valuation-driven approaches, scenario analysis simplifies the assessment process, but in doing so runs the risk of fostering complacency as organizations, satisfied that they have addressed the specified scenario risks, are led into a potentially false sense of security.

1.5.4 Best Practices

The establishment of best practices is a common engineering response to the problem of standardization where subjective judgments may cause variations in design and implementation. The idea that computer security, which suffers from inadequate data and a reliance on subjective expert assessments, could also benefit from a best-practices approach should not be surprising. The basic premise is that conformance to a set of best practices will ensure protection against negligence-based liability lawsuits in the event of a security breach. Proponents of this approach would define best practices as the policies and safeguards of a majority of industry participants.²⁰

This alternative approach to risk management is the least analysis-intensive of all the methodologies examined. The need for assessing threats, generating scenarios, analyzing consequences, or collecting security-related data is circumvented by the industry best practices. A set of common security practices can be found in a number of publicly

¹⁹ For more information on Sun Microsystems' computer security risk-management process, see Timothy J. Townsend, *Security Adequacy Review Process and Technology*, Technical White Paper (Palo Alto, CA: Sun Microsystems, 1998). For a more detailed description of JP Morgan's risk management process, see The Economist Intelligence Unit, *Managing Business Risks in the Information Age* (New York: The Economist Intelligence Unit, Ltd., 1998).

²⁰ See Charles C. Wood, *Best Practices in Internet Commerce Security* (Sausalito, CA: Baseline Software, Inc., 1998).

available documents, such as British Standard 7799.²¹ Provided that others in the industry follow suit, a set of best practices could be established and, as some contend, a powerful defense erected around the industry against potential liability lawsuits. The viability of this approach depends on the reasonableness of best practices compliance costs and the efficacy of the practices themselves. Provided that the costs of conformance are proportionate to the actual risks and that the security measures are effective, the best practices approach is likely to succeed.

Like the other simplifying approaches, best practices may not be a satisfactory long-term solution to the problem. The refinement and advancement process of best practices could easily become uncoupled from the risks they are intended to address. This detachment could result from a “herd mentality” that compels an organization to follow a practice, regardless of its applicability to the actual risks, merely because a majority of the industry has adopted it. The uncoupling could also be the product of a runaway game between industry participants wherein every organization strives to be “above average” in its practices. The result in either case would be security practices unjustified by the actual risks. Finally, best practices de-emphasizes the need for data collection because risk management is not linked in any way to a quantitative analysis. Thus, not even a retrospective analysis of the standard’s effectiveness, cost-benefit trade-offs, or even basic applicability to threats can be performed. Like the other approaches, however, best practices might be an effective means of jump-starting the process of improving security. Many members of SRI’s International Information Integrity Institute (I4) have, in frustration, forsaken formal frameworks of risk assessment and adopted the best practices approach instead.²²

1.6 Underlying Forces for Change

The four new approaches to managing risks are, in general, short-term solutions to the problem. Although they have been relatively successful in addressing organizational acceptance and deployment issues, they have left largely untouched the technological and informational challenges of the previous-generation methodologies. Their lack of cost justification, inability to forecast future trends, and disregard for safeguard efficacy measurement will impel organizations to seek a more satisfactory, quantitative framework for managing computer security risks.

The challenges facing computer security risk management are not unique. Financial markets, the insurance industry, and others have dealt with risks in the face of uncertainty, lack of adequate statistics, and technological changes. As Bernstein’s hallmark of the modern era comes to computer security, risk will be measured, distributed, mitigated, and accepted. Data will be collected, and standards will be established. Three driving forces will motivate and shape the emergence of a new quantitative framework: insurance needs, liability exposure, and market competition.

Insurance companies, sensing an enormous market opportunity, are already testing the waters of computer-security-related products. Safeware, American Insurance Group, and others now offer a variety of policies with coverage ranging from hardware replacement to full information-asset protection. As claims are made against these policies, the industry will begin building actuarial tables upon which it will base premiums. Inevitably, classifications of security postures will develop, and an organization’s security rating will dictate the coverage it can obtain and the premium it must pay. These advances are

²¹ British Standard 7799 can be purchased from the BSI Online store at <<http://www.bsi.org.uk/index.xhtml>>.

²² For more information on the best practices concept, see Donn B. Parker, *Fighting Computer Crime: A New Framework for Protecting Information* (New York: John Wiley & Sons, Inc., 1998).

inescapable and not dependent upon the cooperation of organizations in an information-sharing regime. At present, one of the major obstacles preventing the production of actuarial tables is the widespread reticence to share the little data that has been collected. However, with the widening role of insurance, information sharing must, at some level, take place for claims to be filed and compensation for losses to occur. In this way, metrics of safeguard efficacy, incident rates, and consequences will emerge, and a new quantitative risk framework will begin to take shape.

Avoidance of legal liability is commonly cited as a reason for improving information security. Organizations often find themselves in possession of confidential or proprietary information belonging to third parties with whom they have no formal agreement. If that information should be somehow lost or stolen, thus causing injury to its original owner, the organization may be liable. Under the strictures of tort law, the somewhat vague standard of the “reasonable man” is used to judge liability of negligence.²³ Advocates of the best practices approach argue that compliance to an industry standard will, in and of itself, protect an organization against liability lawsuits. This contention is based on a strain of legal reasoning that can be traced back to an 1890 Pennsylvania railroad liability case. The relevant ruling in the case states that “No jury can be permitted to say that the usual and ordinary way, commonly adopted by those in the same business, is a negligent way for which liability shall be imposed.”²⁴

Since 1890, however, tort law has evolved and with it the standard by which negligence is decided. In 1932 Judge Learned Hand directly addressed the issue of whether an industry custom constitutes a sufficient defense against charges of negligence in the case of T. J. Hooper:

There are, no doubt, cases where courts seem to make the general practice of the calling the standard of proper diligence; we have indeed given some currency to the notion ourselves. Indeed in most cases reasonable prudence is in fact common prudence; but strictly it is never its measure; a whole calling may have unduly lagged in the adoption of new and available devices. It never may set its own tests, however persuasive be its usages. Courts must in the end say what is required; there are precautions so imperative that even their universal disregard will not excuse their omission.²⁵

Fifteen years later, in *United States vs. Carroll Towing Company*, Judge Hand once again returned to the question of liability and negligence. In that case, he articulated a formula that has gone on to become one of the defining tests of negligence.²⁶

²³ Lawrence M. Friedman, *A History of American Law*, 2nd edition (New York: Simon & Schuster, 1985), p. 468.

²⁴ *Titus vs. Bradford, B. & K. R. Co.*, 20 A. 517 (Pa. 1890) as quoted in Richard A. Epstein, *Cases and Materials on Torts*, 6th edition (Boston: Little, Brown & Co., 1995), p. 213.

²⁵ T.J. Hooper, 60 F.2d 737 (2d Cir. 1932) as quoted in Richard A. Epstein, *Cases and Materials on Torts*, 6th edition (Boston: Little, Brown & Co., 1995), p. 216.

²⁶ See *United States vs. Carroll Towing Company*, 159 F.2d 169, 173 (2d Cir. 1947).

Let :

P = Probability of injurious event

L = Gravity of the resulting injury

B = Burden, or cost, of adequate precautions

Then,

Injurer is negligent if and only if $B < PL$

Thus, the costs of avoiding an accident and the expected cost of the accident “must be compared at the margin, by measuring the costs and benefits of small increments in safety and stopping investing in more safety at the point where another dollar spent would yield a dollar or less in added safety.”²⁷

At first glance, this standard might appear to contradict the best-practices advocates who argue that compliance with an industry standard will protect an individual firm from liability actions. Although Judge Hand’s cost-benefit test has its detractors, it “has received overwhelming acceptance in the courts.”²⁸ A reconciliation of sorts can be found in the assumption that the best-practices standard would be set, ostensibly, by “organizations that take prudent care to protect their information.”²⁹ If one assumes that these organizations are rational acting, profit-maximizing entities, then their self-interest will lead them to implement cost-effective information security policies, and, in such a case, to paraphrase Judge Hand, common prudence will in fact be reasonable prudence. The final arbiter, however, will likely be Judge Hand’s cost-benefit standard, thus creating an incentive for organizations to collect the necessary data that will enable them to justify their information security policies with credible assessments of risk.

Competition and market forces are probably the last great engine of change that will force companies to protect their information assets efficiently. Regardless of the risk-management strategy pursued, whether it be ALE-based assessment, scenario analysis, best practices, or some other, the marketplace will ultimately govern whether that strategy was an efficient use of resources. Those companies that secure their assets cost-effectively will gain a competitive advantage over those that do not. Thus, businesses that over-protect will have spent too much on security, and those that under-protect will suffer greater losses as a result of ensuing security breaches.

Insurance, liability, and competition are underlying forces that will push computer security risk management away from the non-quantitative, second-generation approaches and back toward a quantitative framework of risk assessment and management similar to the first-generation ALE-based methodologies. They will enable, and in some cases insist upon, the quantification of risks, the measurement of safeguard efficacy, and the analysis of costs and benefits. Thus, a new approach is needed, tempered by the lessons of the past and capable of adapting to the future demands that these forces will soon thrust upon it.

²⁷ Richard A. Posner, *Economic Analysis of Law*, 4th edition (Boston: Little, Brown & Co., 1992), p. 164.

²⁸ Richard A. Epstein, *Cases and Materials on Torts*, 6th edition (Boston: Little, Brown & Co., 1995), p. 218.

²⁹ Donn B. Parker, *op. cit.*, p. 284.

Chapter 2 Risk Modeling and Analysis

“The last truth is that there is no magic,” said Magician Pug.
—Raymond E. Feist, *Prince of the Blood*³⁰

Modeling and simulation can, at times, take on an almost mystical quality. Good models are able to predict future system behavior reliably, and foretelling the future has always commanded a certain reverence. Unfortunately, that reverence sometimes leads to an ascription of greater significance and capability than is warranted. In the end, models are only reflections of human thought and observation. They do not solve problems on their own and require quality input information to perform their “magical” clairvoyance. Modeling in the absence of adequate supporting data is of little use to anyone.

This dependence on quality data has been, in general, underappreciated or entirely ignored by past computer security risk modelers. The ALE-based risk models simply assumed that the millions of scenario assessments could be and would be readily performed. The second-generation approaches, while at least recognizing the dearth of quality data, neglected to propose or to encourage any policies that would alleviate this problem, opting instead to evade the issue entirely.

Nevertheless, the modeling of risk is an important endeavor. The apparent lack of reliable, relevant data does not excuse organizations from managing their computer security risks. Decisions about security resource allocations must and will be made. Therefore, the duty of a risk model is to inform that resource decision by providing the best possible estimations of risk exposure, policy option efficacy, and cost-benefit analysis. To date, risk models have been preoccupied by a futile quest for comprehensive coverage of all possible hazards. The ensuing controversies over vulnerabilities, probabilities, frequencies, valuations, and other details have hindered and, in many cases, completely derailed efforts to inform risk management decisions. In this chapter, I will propose a decision analytic framework for managing risks that addresses many past problems by shifting modeling focus away from the details of computer security interactions and placing it squarely on the risk-management decision.

2.1 Risk Modeling As a Decision-Driven Activity

The National Research Council, recognizing similar problems in other areas where risk management is a matter of public policy, recommended a new focus for what it termed “risk characterization” in *Understanding Risk: Informing Decisions in a Democratic Society*.³¹ According to the report, risk characterization, or a summary of technical analysis results for use by a decision maker, should necessarily be a decision-driven activity, directed toward informing choices and solving problems. Risk characterization, the report goes on to say, should emerge from

an analytic-deliberative process . . . [whose] success depends critically on systematic analysis that is appropriate to the problem, responds to the needs of the interested and affected parties, and treats uncertainties of importance to the decision problem in a comprehensible way. Success also

³⁰ Raymond E. Feist, *Prince of the Blood* (New York: Doubleday, 1989), p. 67.

³¹In this context, risk characterization and risk modeling are synonymous.

depends on deliberations that formulate the decision problem, guide analysis to improve decision participants' understanding, seek the meaning of analytic findings and uncertainties, and improve the ability of interested and affected parties to participate effectively in the risk decision process.³²

Although the National Research Council's report concentrates specifically on a public-policy process, its lessons are nevertheless instructive for private organizations. Casting risk characterization as a decision-driven activity recognizes the fact that some policy will be inevitably chosen, even if that policy is to do nothing. Implicit in this decision are assessments of key variables and determinations of value trade-offs, and the proper role of risk modeling is to make those assessments and determinations explicit for decision makers. This process allows decision makers to better understand the ramifications of their choices, to weigh their beliefs within the decision framework, and to be cognizant of the underlying assumptions upon which their decision will be based.

Decision-driven analyses of complex problems involving uncertainty, incomplete data, and large investments are not unknown to private industry. The business literature is replete with books and articles describing how companies can better manage their research and development portfolios, product transitions, inventory maintenance, and a myriad of other problems common to businesses.³³ In this way, the integrated business risk approach made an important contribution by recognizing and emphasizing the placement of computer security risks among other operational risks faced by businesses. In this chapter, I will present a quantitative decision analysis framework for assessing and managing computer security risks as a candidate decision-driven analytic modeling process.

2.2 Decision Modeling

The application of statistical decision theory to management problems traces its roots to the seminal work of Raiffa and Schlaifer in 1961³⁴ with considerable refinement by Howard in 1966.³⁵ The term "decision analysis" was coined by Howard to refer specifically to the formal procedure for analyzing decision problems outlined in his article and subsequent research. At its core, decision analysis is a reductionist modeling approach that dissects decision problems into constituent parts: decisions to be made, uncertainties that make decisions difficult, and preferences used to value outcomes.

Decision analysis offers several key advantages that recommend it well to the problem of computer security risk management. First, as its name implies, it is necessarily a decision-driven modeling technique. Second, its incorporation of probability theory provides it tools to capture, clarify, and convey uncertainty and the implications of uncertainty. Third, and probably most important, decision analysis utilizes influence diagrams as a common graphical language for encapsulating and communicating the collective knowledge of an organization, thus facilitating consensus-building.

Influence diagramming, by mapping the relationships between key variables, can be a powerful tool for both communicating and analyzing the underlying factors that influence

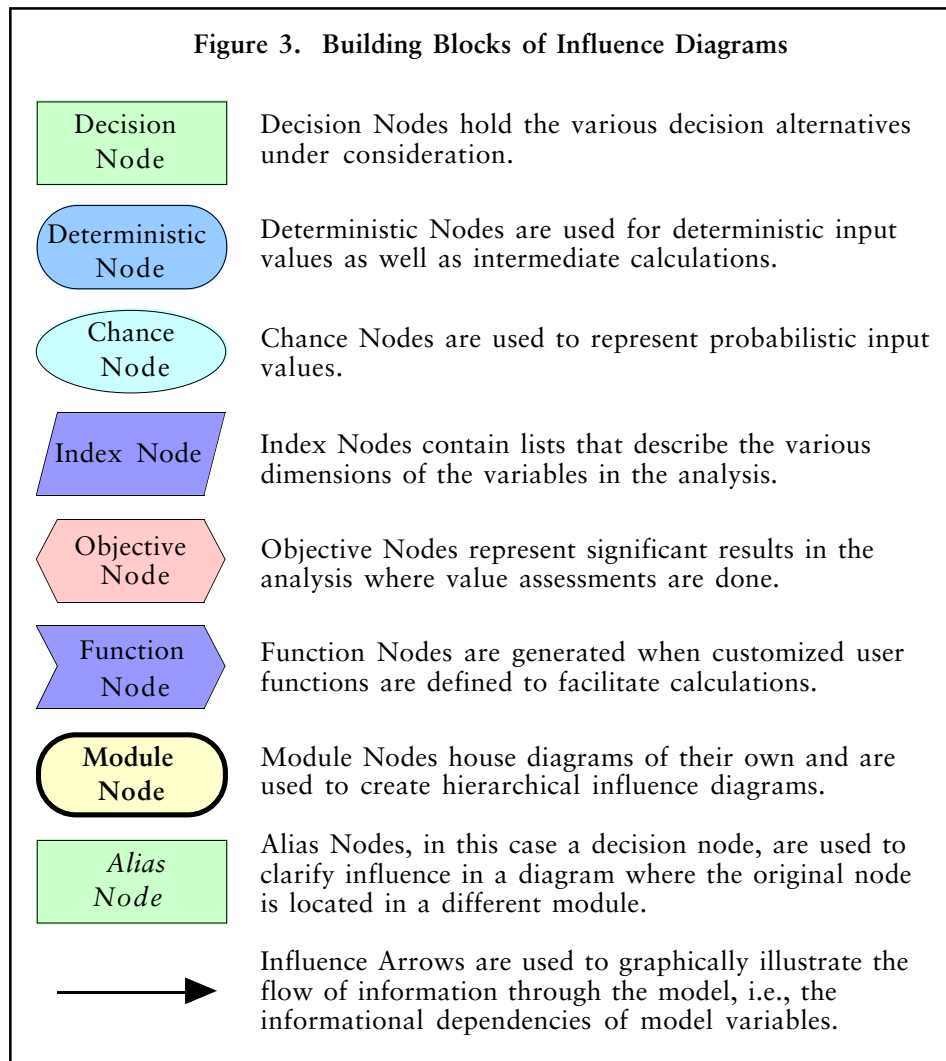
³² Stern and Fineberg, *op. cit.*, p. 3.

³³ For example, see *Harvard Business Review on Managing Uncertainty* (Boston: Harvard Business School Press, 1999); Robert G. Cooper, Scott J. Edgett, and Elko J. Kleinschmidt, *Portfolio Management for New Products* (Reading, MA: Addison-Wesley, 1998); or David Matheson and Jim Matheson, *The Smart Organization: Creating Value through Strategic R&D* (Boston: Harvard Business School Press, 1998).

³⁴ Howard Raiffa and Robert Schlaifer, *Applied Statistical Decision Theory* (Boston: Harvard University, 1961).

³⁵ Ronald A. Howard, "Decision Analysis: Applied Decision Theory," *Proceedings of the Fourth International Conference on Operational Research*, David B. Hertz and Jacques Melese, editors (New York: Wiley-Interscience, 1966), pp. 55–71.

decision making.³⁶ The diagrams are composed of nodes, representing variables, and arrows, representing influence between variables. By convention, the shape of a node dictates the type of variable it represents and the kind of influence represented by arrows originating from or pointing to it; see Figure 3 for descriptions of the different node types.



Arrows are used to demonstrate relationships between variables in a decision model and thus the flow of information through a model.³⁷ Specifically, an arrow directed into a deterministic node or objective node indicates that the destination node is a function of the origin node. An arrow into a chance node denotes probabilistic dependence, meaning that

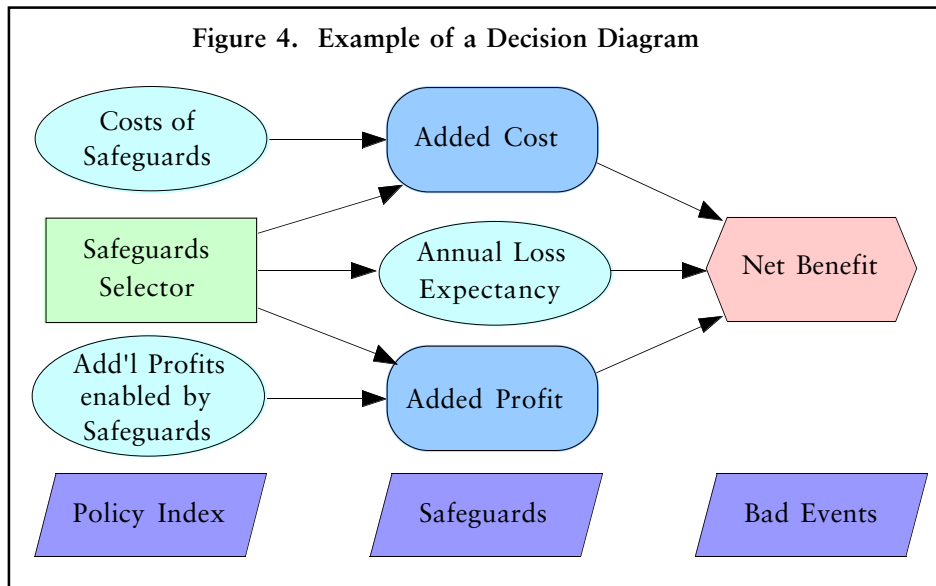
³⁶ Ronald A. Howard and James E. Matheson, "Influence Diagrams," *Readings in the Principles and Applications of Decision Analysis*, Vol. 2, Ronald A. Howard and James E. Matheson, editors (Menlo Park, CA: Navigant Consulting, Inc., 1983), p. 719.

³⁷ The diagramming conventions in this paper reflect those of the Analytica™ Decision Modeling Environment. Analytica, developed under the original name of DEMOS™ at Carnegie Mellon University, is one of the more powerful commercially available, influence-diagram-based, decision analysis software packages. A free evaluation version may be downloaded from Lumina Decision Systems' web site at <<http://www.lumina.com>>. The risk model described in this thesis may be viewed, analyzed, and evaluated with the free evaluation version. See the Appendix for more information.

the chance node's probability distribution is conditioned on the origin node's values. An arrow into a decision node represents informational influence, indicating that the origin node's value will be known at the time of the decision and may thus affect the decision.³⁸

The construction of a decision-oriented influence diagram, or decision diagram,³⁹ is in itself a valuable learning and consensus-building exercise. The first step is to define the decision and all alternative courses of action that may be taken in making that decision. This variable is usually represented with a rectangular decision node and placed on the left side of the diagram. Once the decision is articulated, the values and objectives must be codified, and metrics must be established. Ideally, the metrics can be functionally related, thus allowing a single objective node to be defined and placed on the right side of the diagram. Next, working from right to left, from objective to decision nodes, chance and deterministic nodes should be inserted to represent uncertainties, available input data, and intermediate calculations. An example decision diagram for a computer security risk management decision is given in Figure 4.

In Figure 4, the primary decision to be made is the selection of security safeguards. Rather than compare all safeguards on an individual basis, this model allows the user to group safeguards together into baskets of safeguards, or policies, and to make comparisons between policies. These policy names are listed in the *Policy Index* variable. Because index nodes influence nearly every other node in the diagram, influence arrows originating from them are often concealed to enhance diagram clarity. The *Net Benefit* objective node on the right contains a policy cost-benefit trade-off calculation. This trade-off depends upon three factors: added implementation costs of safeguards, additional profits expected from



new opportunities, and annual loss expectancy. As the arrows emanating from the decision node indicate, these factors will vary according to the safeguards selected under each policy. The three index nodes across the bottom of the diagram contain lists that describe the dimensions of the analysis. For example, the *Policy Index*, as was mentioned

³⁸ M. Granger Morgan and Max Henrion, *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*, 2nd edition (New York: Cambridge University Press, forthcoming), p. 262.

³⁹ Strictly speaking, influence diagrams do not necessarily involve decisions. Knowledge maps, for example, are influence diagrams and may not contain any decision nodes.

before, contains the names of the policies under consideration; the *Safeguards* index node lists the safeguards that can be selected; and the *Bad Events* index node contains the security compromises under evaluation.

As with any modeling effort, a balance must be struck between model fidelity and tractability. As it has been applied in professional practice, decision analysis tends to approach this balance from the side of model tractability. Through an iterative process of progressive refinement, decision models evolve, becoming more complex as model analysis, data availability, and data relevance indicate a need for and capability of providing greater detail. Any one of the variables shown in the diagram in Figure 4 could be replaced with a module, representing an entire influence diagram composed of other, more elementary quantities. See Figure 5 for an example where Annual Loss Expectancy has been thusly transformed.

With hierarchical influence diagrams such as these, very large, complex systems may be modeled without losing the elegance and communicative benefits of influence diagrams. As the diagram shows, annual loss expectancy depends upon the expected annual frequency of bad events, or security compromises, and upon the expected consequences of those bad events. These intermediate calculations are a function of a deterministic reduction factor, a probabilistic estimate of an initial value, and the safeguards selected for each policy.⁴⁰

2.3 Computer Security Risk Model Description

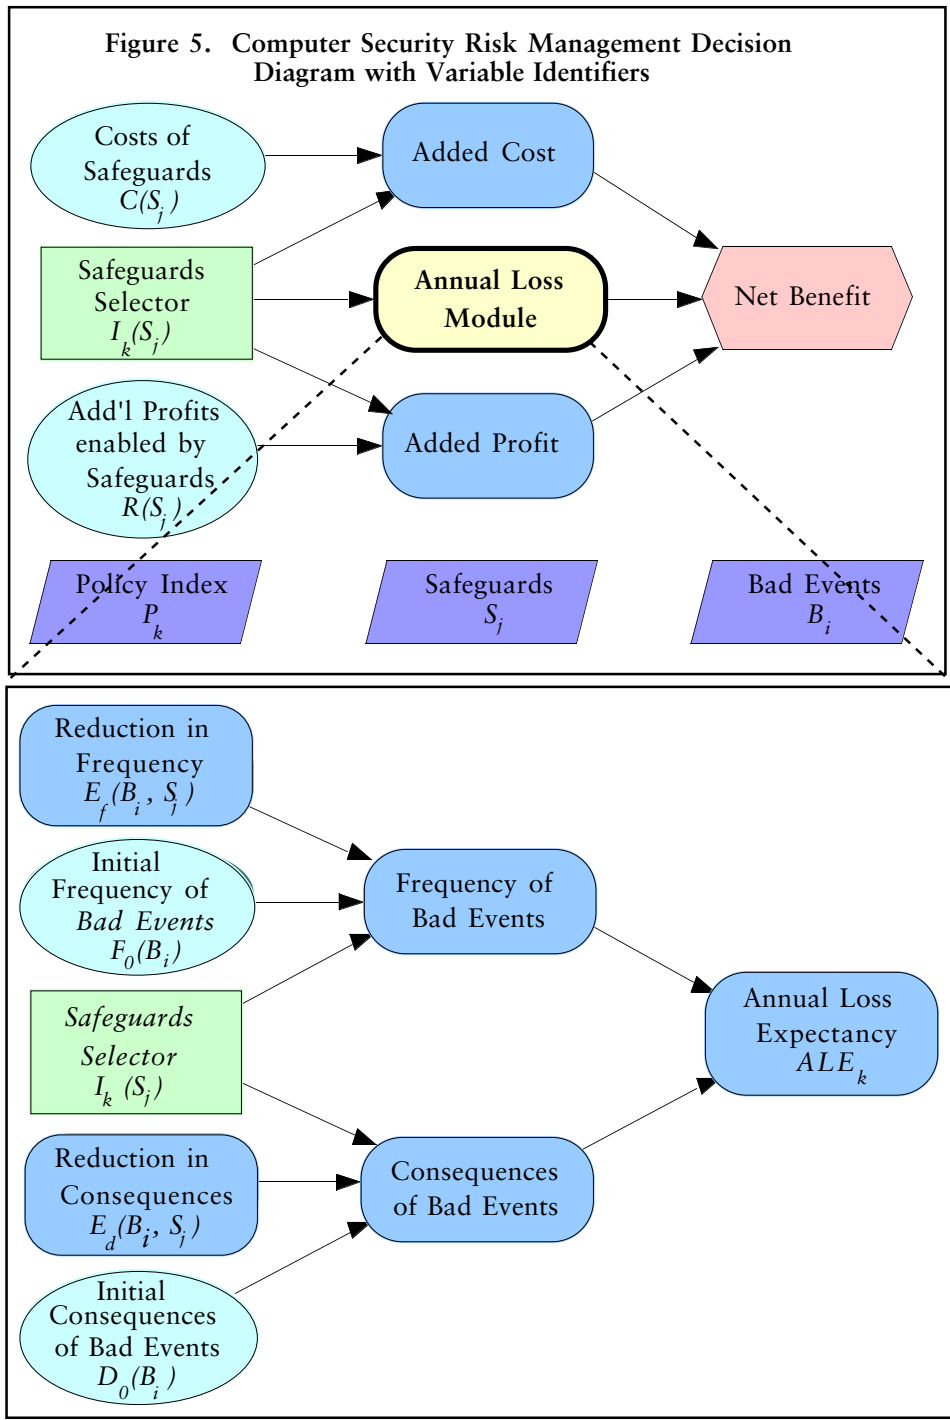
Figure 5, in addition to providing a rudimentary illustration of hierarchical decision diagrams, also presents a fairly concise summary of the essential quantities that must be assessed, either implicitly or explicitly, whenever a computer security risk management decision is made. A more thorough exploration of this diagram would serve the dual purpose of demonstrating the model construction process and elucidating these fundamental variables.

For mathematical clarity and facility, the following variable identifiers will be used.

B_i	=	Bad event i where $i = \{1, 2, 3, \dots, n\}$ For example, data theft, service outage, employee theft
S_j	=	Safeguard j where $j = \{1, 2, 3, \dots, m\}$ For example, awareness program, firewalls, encryption software
P_k	=	Policy k where $k = \{0, 1, 2, 3, \dots, l\}$ For example, status quo, incremental change, major improvement By convention, $k = 0$ represents the status quo.
$R(S_j)$	=	New profits enabled by adoption of safeguard S_j
$I_k(S_j)$	=	Binary function indicating if safeguard S_j is included in policy P_k
$F_0(B_i)$	=	Initial estimate of the frequency of bad event B_i
$D_0(B_i)$	=	Initial estimate of the consequences of, or damage from, the occurrence of bad event B_i
$E_f(B_i, S_j)$	=	Fractional reduction in frequency of occurrence of bad event B_i as a result of implementing safeguard S_j
$E_d(B_i, S_j)$	=	Fractional reduction in consequences resulting from the bad event B_i as a result of implementing safeguard S_j
$C(S_j)$	=	Cost of implementing safeguard S_j
ALE_k	=	Annual loss expectancy under policy P_k

⁴⁰ For more information about decision modeling in Analytica, see M. Granger Morgan and Max Henrion, *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*, 2nd edition (New York: Cambridge University Press, forthcoming), pp. 257–290.

Figure 5. Computer Security Risk Management Decision Diagram with Variable Identifiers



Beginning with the decision, the *Safeguards Selector* is essentially a matrix of ones and zeroes indicating whether a particular safeguard is included in a policy or not. Thus, the decision problem is framed as one of comparing alternative policies, i.e., baskets of safeguards, to determine which policy is best. The criteria for that determination is sometimes called a utility function, representing the preferences of the individual or group responsible for making the decision. The utility function for the diagram in Figure 5 is captured in the *Net Benefit* objective node, the equation of which is given below.

$$Net\ Benefit_k = Benefit_k - Added\ Cost_k + Added\ Profit_k \quad \forall k = \{1,2,3,\dots,l\} \quad (3)$$

Note: Because the decision is framed as a choice between competing policies, a net benefit calculation must be done for each policy P_k .

Coefficients or other functional operations may be introduced into the utility function to express risk tolerance, time-value of money, or other preferences. Thus, utility values often have meaningless units and are useful only in a comparative framework. For this analysis, however, the framework has been kept simple, omitting these additional factors, assuming risk neutrality, and keeping real dollars for units.

Calculating the benefits of a security policy can very quickly become a complex and controversial exercise. When people begin considering intangible concepts, such as peace of mind, reputation, and public trust, they tend to place a premium on the value of security beyond the calculated savings that result from avoiding security breaches. However, the weight given to such attributes may vary significantly in both magnitude and manner, depending upon individual or group preferences. Because the first cycle in a decision analysis should be as tractable as possible, the intangible values are best left for later iterations. Thus, the differences in annual loss expectancies for the different policies, relative to the status quo ($k = 0$), provide a conservative estimate of their expected benefits and hence the added value of security.

$$Benefit_k = ALE_0 - ALE_k \quad \forall k = \{1,2,3,\dots,l\} \quad (4)$$

As discussed in the previous chapter, ALE is the product of an incident's annual frequency times its total losses. Actual values for these component quantities will be examined more closely in the next chapter, but losses, in general, include all additional costs incurred as a direct result of an incident. For example, legal fees from an ensuing liability lawsuit, regulatory violation penalties for financial information prematurely disclosed, lost earnings from delayed product development, and lost earnings due to a diminished reputation could all be included in the total loss estimate for an incident. Because insurance premiums depend upon ALE, ALE is likely to become a standard metric. Another of its advantages rests in the rigor and additive nature of its calculations. Unlike measures of intangible values which are subjective and often controversial, the ALE gives real losses and savings under different security policies. Thus, management decisions based on ALE can be cost-justified, even if all possible risks are not considered, because the considered risks alone will warrant the actions taken.⁴¹ The primary drawback of using ALE as an exclusive measure of risk, as mentioned before, rests in its inability to differentiate low-consequence, high-frequency events from catastrophic-consequence, low-frequency events. For many organizations, the former is quite manageable as an incremental cost while the latter would be disastrous. Thus, if an analysis included the

⁴¹ The same cannot be said for comparative risk methods in which scenarios are compared and ranked by severity. In such procedures, an unconsidered risk could invalidate or substantially alter an analysis and its conclusions and recommendations.

latter type of event, the analyst would have to preface the risk assessment findings with an explanation of this failing in the ALE metric. The equation for the *Annual Loss Expectancy* variable is given below.

$$ALE_k = \sum_{i=1}^n \left\{ F_0(B_i) D_0(B_i) \prod_{j=1}^m \left[(1 - E_f(B_i, S_j) I_k(S_j)) (1 - E_d(B_i, S_j) I_k(S_j)) \right] \right\} \quad (5)$$

The concept of safeguard efficacy is needed to weigh the costs and benefits of different security policies. First-generation methodologies' implicit treatment of efficacy manifested itself in the multiple assessments performed under different policy assumptions. In Figure 5, however, efficacy is explicitly modeled as fractional reduction factors in the frequency and consequences of bad events, E_f and E_d , respectively. These reduction factors are relative to the status quo. For example, if $E_f(B_i, S_j) = 0.4$ for some i, j , then the implementation of safeguard S_j would reduce the frequency of bad event B_i to 60 percent of its initial baseline value. When several safeguards are implemented, the reduction factors are applied in succession. For example, if two safeguards, having reduction factors of $E_f(B_i, S_1) = 0.4$ and $E_f(B_i, S_2) = 0.2$, were implemented, then the annual frequency of bad event B_i would be reduced to 48 percent of its status quo value. This simplified approach represents a middle ground in the spectrum of combinatorial algorithms. Some algorithms, such as simple addition of the factors, assume that the sets of incidents prevented by the different safeguards are totally non-overlapping, while others, such as application of only the largest reduction factor, assume that the sets are completely overlapping. The introduction of weighing factors to account for the different degrees of overlap might bring about a more precise overall reduction factor, but at the expense of considerably more assessment.

One of the most important security value propositions to emerge in recent years is security as an enabler of new business opportunities. The prospects of leveraging the Internet for increased sales, closer ties with business partners, expanded internal communications capabilities, and a host of other business-enhancing activities have been a driving force for better computer security. On a cost-benefit basis, these new opportunities could more than justify the added expense of security. Capturing this effect in a security risk model is no easy task. Modeling a new business venture is a complicated exercise in its own right. In a first iteration of the security risk model, a simple estimation of future profits with appropriate uncertainty bounds is sufficient.⁴² If analysis later shows that these new earnings are material to the decision, then more of the new business model could be incorporated in subsequent iterations. The basic equation for *Added Profits* is given below.

$$Added Profit_k = \sum_{j=1}^m R(S_j) I_k(S_j) \quad \forall k = \{1, 2, 3, \dots, J\} \quad (6)$$

Of the key variables, cost is probably the most readily quantifiable. The accounting of safeguard costs typically includes initial investments in equipment and software, salaries of workers doing implementation, and any maintenance costs. Not all costs, however, are so easily assessed. Additional computer security measures can affect employee morale, consume scarce computer and network resources, and hamper ongoing projects, all of which can lead to a drop in productivity. Like the simplifications for *Benefits* and *Added*

⁴² The estimation of profits should be appropriately discounted for risk and time, especially if they will not materialize within the model's time horizon.

Profits, the difficult-to-quantify costs should be left to a subsequent iteration in the analysis cycle. If the decision should turn on costs, then these other factors could be considered. The equation for *Added Cost* is given below, as is the full equation for *Net Benefit*.

$$\text{Added Cost}_k = \sum_{j=1}^m C(S_j)I_k(S_j) \quad (7)$$

$$\begin{aligned} \text{Net Benefit}_k = & \sum_{i=1}^n \left\{ F_0(B_i)D_0(B_i) \left[1 - \prod_{j=1}^m \left[(1 - E_f(B_i, S_j)I_k(S_j))(1 - E_d(B_i, S_j)I_k(S_j)) \right] \right] \right\} \\ & - \sum_{j=1}^m C(S_j)I_k(S_j) + \sum_{j=1}^m R(S_j)I_k(S_j) \end{aligned} \quad (8)$$

2.4 Analysis Techniques

Once the influence diagram and accompanying assessments are satisfactorily complete, the analysis phase ensues. While the diagramming process concentrates institutional knowledge and forges a concurrence of opinion, the analysis phase elucidates critical insights about that knowledge and directs future modeling efforts. The suite of analytical tools for decision model analysis includes nominal range sensitivity analysis, parametric sensitivity analysis, stochastic analysis, and value of information calculations.

Nominal range sensitivity analysis developed as a filtering technique for identifying the key variables that have the greatest influence on a decision. Assessments are made of upper and lower bounds for all uncertainties, usually corresponding to the 10 percent and 90 percent fractiles.⁴³ The objective function is then computed for each upper and lower bound, individually, for a given variable with all other variables kept at their nominal, or “best guess,” values. Generally, those variables found to greatly affect the objective function are targeted for further attention in subsequent analysis and potentially more detailed characterization in other iterations of the modeling process.

Parametric sensitivity analysis, as its name implies, parameterizes a given variable and examines how the optimal decision changes throughout a range of values. Of key interest are the “cross-over points,” where the best decision changes from one decision alternative to another. This sensitivity analysis gives us insight into the robustness of the chosen alternative. If the best decision alternative were to change with small variations in a particular variable, then that variable should be targeted for more detailed modeling and information gathering.

After sensitivity analysis has identified the important variables, stochastic analysis examines the role of uncertainty in the decision.⁴⁴ The rules of stochastic dominance provide techniques for ranking decision alternatives based on their full probabilistic distributions and basic decision-maker preferences. First, second, and third degree stochastic dominance are successively finer criteria for judging superiority among decision alternatives. The following set of equations gives the rules for stochastic dominance.

⁴³ The p fractile, X_p , of a distribution is a value such that there is a probability p that the actual value of the random variable will be less than that value: $P[X \leq X_p] \equiv p$

⁴⁴ Much of this discussion on stochastic dominance is drawn from Haim Levy, *Stochastic Dominance: Investment Decision Making under Uncertainty* (Boston: Kluwer Academic Publishers, 1998), pp. 41–111.

Given :

$A1[X] = P[X \leq x]$ for Alternative 1

$A2[X] = P[X \leq x]$ for Alternative 2

$U =$ Underlying Utility Function

First Degree Stochastic Dominance :

$A2[X]$ dominates $A1[X]$ if and only if $A2[X] \leq A1[X] \forall X$

with strict inequality for at least one value of X .

Assumption : U is non - decreasing, i.e., $U' \geq 0$.

Second Degree Stochastic Dominance :

$A2[X]$ dominates $A1[X]$ if and only if $\int_a^X A2[t]dt \leq \int_a^X A1[t]dt \forall X$

with strict inequality for at least one $X_0 \in [a, b]$.

Assumptions : U is non - decreasing, $U' \geq 0$, and U exhibits risk aversion, $U'' \leq 0$.

Third Degree Stochastic Dominance :

$A2[X]$ dominates $A1[X]$ if and only if

$$(1) \int_a^X \int_a^z A2[t]dtdz \leq \int_a^X \int_a^z A1[t]dtdz \forall X \text{ and}$$

$$(2) \int_a^b A2[X]dX \leq \int_a^b A1[X]dX$$

with strict inequality for at least one X in condition 1 or strict inequality in condition 2.

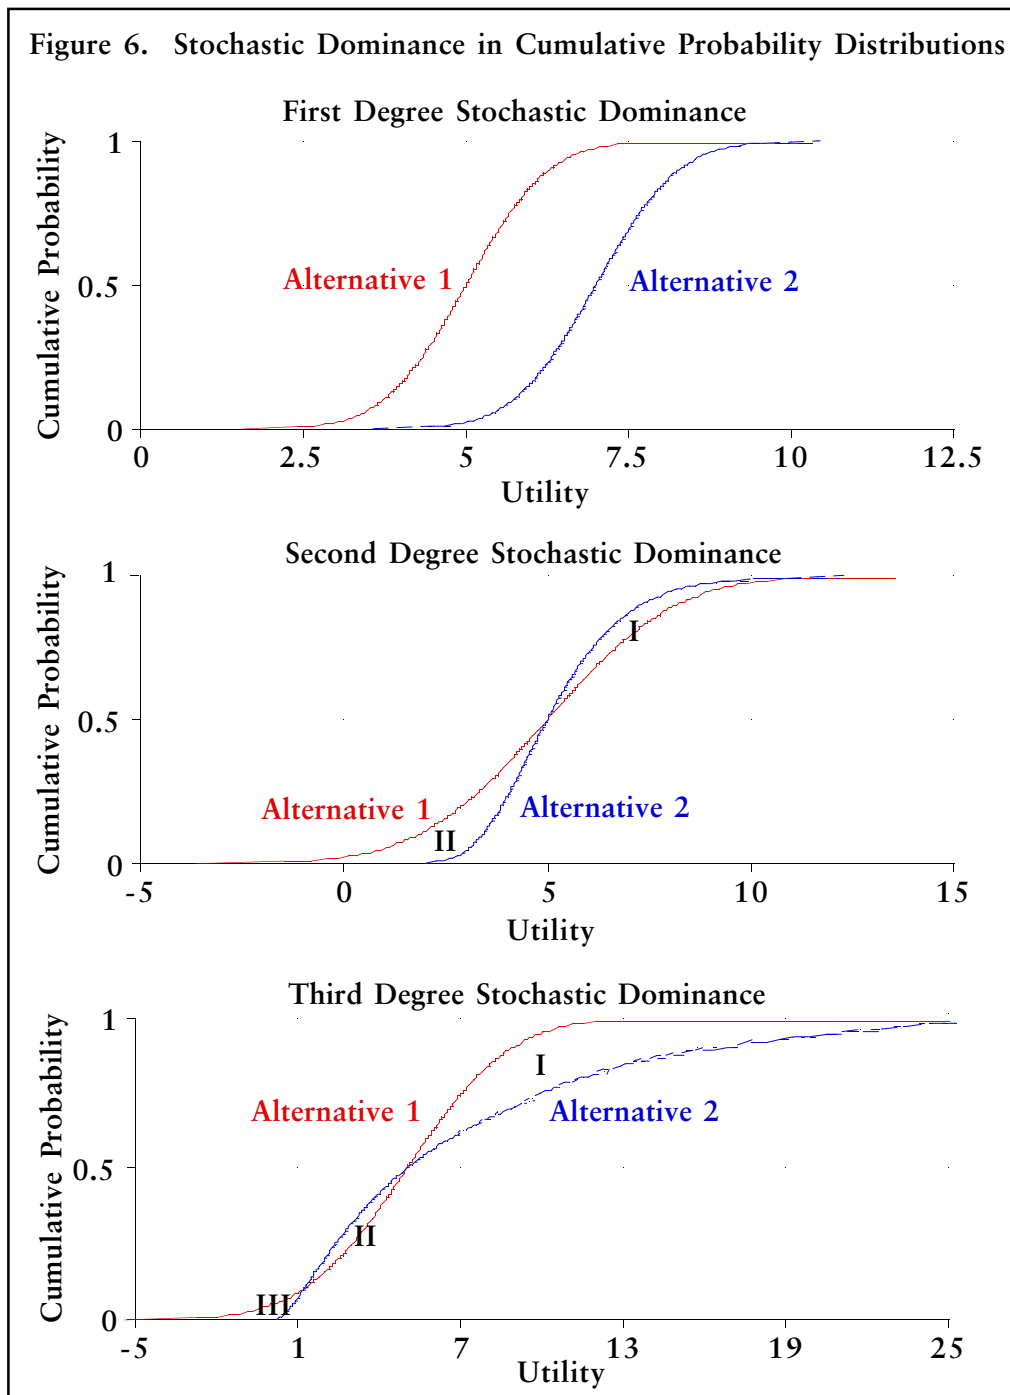
Assumptions : U is non - decreasing, $U' \geq 0$; U exhibits risk aversion, $U'' \leq 0$, and

U exhibits positive skewness, $U''' \geq 0$.

Stochastic dominance is probably easiest seen by looking at the cumulative probability distribution graphs of competing decision alternatives, as in Figure 6. In the first graph, first-degree stochastic dominance (FSD) can be concluded by inspection because the dominant distribution curve of Alternative 2 is fully to the right of the curve for Alternative 1. One way to understand this result is to consider that for any utility value the probability of achieving that utility or higher is always greater with Alternative 2. Thus, the only underlying assumption behind FSD is that the utility function is non-decreasing. If the distribution curves were to cross, then FSD would not exist. In such a circumstance, one might look for second-degree stochastic dominance (SSD) by comparing the areas between the curves. As the second graph in Figure 6 shows, Alternative 2 has second-degree stochastic dominance over Alternative 1 because area II is larger than area I. The underlying assumptions for SSD are that the utility function is non-decreasing and exhibits risk aversion.⁴⁵ The intuitive interpretation of SSD is that a decision maker will

⁴⁵ The idea behind risk aversion is that given a choice between two options of the same expected value but different variances, the risk-averse person will choose the option with the smaller variance. Variance is a measure of a distribution's "spread" or uncertainty. The greater the variance, the greater the range of likely values.

Figure 6. Stochastic Dominance in Cumulative Probability Distributions



prefer less risky alternatives, ones with smaller variances, to those of greater risk, ones with larger variances. The third graph in Figure 6 shows an example of third-degree stochastic dominance (TSD). One might be tempted to think that SSD exists here because the sum of areas I and III is clearly larger than area II. However, area II is larger than area I, and, because the definition of SSD requires that the inequality hold for all of X, SSD is not present. Graphically, TSD requires that the sum of areas I and III be greater than area II

II and that area II not be more than three times as large as area I. The underlying assumptions of TSD are that the utility function is non-decreasing, exhibits risk aversion, and demonstrates positive skewness. An intuitive interpretation of TSD is that given alternatives with equal means and variances, a decision maker will prefer the alternative with a lower “downside risk,”⁴⁶ as manifested in a positive skewness.

In the absence of any stochastic dominance, greater modeling effort and information gathering may help to narrow the uncertainty of key variables and bring about a more refined understanding in which dominance can be found. Value-of-information analysis is helpful for determining the resources that should be spent in pursuit of additional information. In general, the idea that information has value is typically an undisputed fact; pegging a precise number to that value, however, is far more contentious. In the context of a decision, this task becomes somewhat more straightforward. To the extent that the discovery of additional information about an uncertainty can change a decision, that information has value. Knowing a variable with certainty changes the expected utility of each policy alternative. By taking the difference between the utility of the decision with new information and the utility of the decision without new information, we can derive a value for that new information.

For example, if $\bar{C}(S_j) =$ Information on the cost of safeguard S_j , then

$$\begin{aligned} \text{Value of Information for } \bar{C}(S_j) = & \hspace{20em} (9) \\ \text{Max}_{k=1}^I \left[\text{Net Benefit} \left(ALE_k, R(S_j), C(S_{i \neq j}), \bar{C}(S_j) \right) \right] & - \text{Max}_{k=1}^I \left[\text{Net Benefit} \left(ALE_k, R(S_j), C(S_j) \right) \right] \end{aligned}$$

Unfortunately, this value is not computable before the information is known. However, an expected value of perfect information (EVPI) can be calculated prior to the information being discovered. EVPI is found by computing the probabilistic expectation of the decision utility with new information less the expectation of the decision utility without new information.

$$\begin{aligned} \text{EVPI for } \bar{C}(S_j) = \sum_{C(S_j)} P \left[C(S_j) = \bar{C}(S_j) \right] \text{Max}_{k=1}^I \left[\text{Net Benefit} \left(ALE_k, R(S_j), C(S_{i \neq j}), \bar{C}(S_j) \right) \right] & \hspace{2em} (10) \\ - \text{Max}_{k=1}^I \left[\text{Net Benefit} \left(ALE_k, R(S_j), C(S_j) \right) \right] & \end{aligned}$$

The EVPI is always positive unless the new information has no bearing on the decision alternative chosen. In this case, the decision maker’s decision is unaffected by the new information, and the information, therefore, has an EVPI of zero.⁴⁷

In 1764, Thomas Bayes’ “Essay towards Solving a Problem in the Doctrine of Chances” introduced a method for blending new information with prior beliefs. The Bayesian framework for incorporating new information should be particularly useful in computer security risk management, where risk models must keep pace with rapid changes in technology and human behavior. By modifying the original “prior” probability distribution of a given uncertainty with a “pre-posterior” distribution, which describes the new information’s predictive accuracy, a new “posterior” distribution for the given uncertainty can be derived. This posterior distribution will then replace the prior

⁴⁶ Downside risk is a business term for the risk of undesirable consequences, such as lower utility.

⁴⁷ For a more rigorous treatment of value of information, see Howard Raiffa and Robert Schlaifer, *Applied Statistical Decision Theory* (Boston: Harvard University, 1961), pp. 87–91.

distribution in all calculations. For example, if new information, $T(S_j)$, were found regarding $C(S_j)$, then

$$\text{Prior distribution of } C(S_j) = P[C(S_j)]$$

$$\text{Pre - Posterior distribution is : } \frac{P[T(S_j) | C(S_j)]}{P[T(S_j)]}$$

$$\text{Posterior distribution is : } P[C(S_j) | T(S_j)] = P[C(S_j)] \frac{P[T(S_j) | C(S_j)]}{P[T(S_j)]}$$

This same framework can also be used to compute the expected value of imperfect information (EVII). Most often, information about an uncertain variable is not perfect. Indeed, perfect information is typically very expensive, if not impossible, to obtain. The following example looks, again, at the cost of safeguard S_j . This time, however, the information, $T(S_j)$, is not perfect. Thus, the true value of $C(S_j)$ may not always equal the indicated value, $\bar{C}(S_j)$.

Let :

$C(S_j)$ = Uncertain cost of safeguard S_j

$T(S_j)$ = Imperfect information about $C(S_j)$

$\bar{C}(S_j)$ = Indicated value of $C(S_j)$ from imperfect information

Then : EVII for $\bar{C}(S_j)$ =

$$\sum_{T(S_j)} P[T(S_j)] \underset{k=1}{\overset{l}{\text{Max}}} \left[\text{Net Benefit}(ALE_k, R(S_j), C(S_j \neq j), \bar{C}(S_j), P[C(S_j) | T(S_j)]) \right] \\ - \underset{k=1}{\overset{l}{\text{Max}}} \left[\text{Net Benefit}(ALE_k, R(S_j), C(S_j)) \right]$$

Thus, the Bayesian framework provides a means for incorporating new information and for valuing imperfect information in a decision context.⁴⁸

2.5 Summary

The decision analysis approach offers several key advantages that address many of the criticisms leveled against past risk models. The approach recognizes that a decision will be made and provides tools for making explicit the roles of expert judgment, past data, and underlying assumptions in the risk assessment. Its top-down, iterative framework prevents the analysis from becoming mired in more detail than is practicable. By starting with a simple problem formulation and using analysis to dictate where greater modeling effort and additional information should be focused, decision modeling is able to keep a tight rein on model complexity. Influence diagramming, with its common language and process

⁴⁸ For more information on decision model analysis techniques, see M. Granger Morgan and Max Henrion, *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*, 1st edition (New York: Cambridge University Press, 1990), pp. 172–220.

for achieving consensus, helps to address deployment issues. Although no modeling technique can completely compensate for a lack of good data, the use of probability distributions to express the limits of knowledge can curtail or even avert controversies over poor data or expert judgments. The data-dependence of this modeling approach grounds the analysis in quantifiable reality and encourages the systematic collection of supporting data to update and improve the risk model. In the next chapter, we will examine the availability, reliability, and relevance of publicly available data for use in risk-modeling activities.

Chapter 3 Data

To measure is to know.
—James C. Maxwell

The collection, analysis, and dissemination of data relevant to computer security risks is widely recognized as indispensable to improving information security.⁴⁹ Information transparency, similar to that which exists for thefts, fires, automobile accidents, and a myriad of other undesirable events, is essential for well-informed risk decisions. Many obstacles have impeded or tainted efforts to initiate and to promote information sharing, including measurement problems, questions of relevance, terminology incompatibilities, competitive pressures, fear of embarrassment, and liability concerns. These impediments must be meaningfully addressed to effect an information-sharing regime adequate for both private risk management activities and national critical infrastructure policy formulation.

3.1 Difficulty with Data

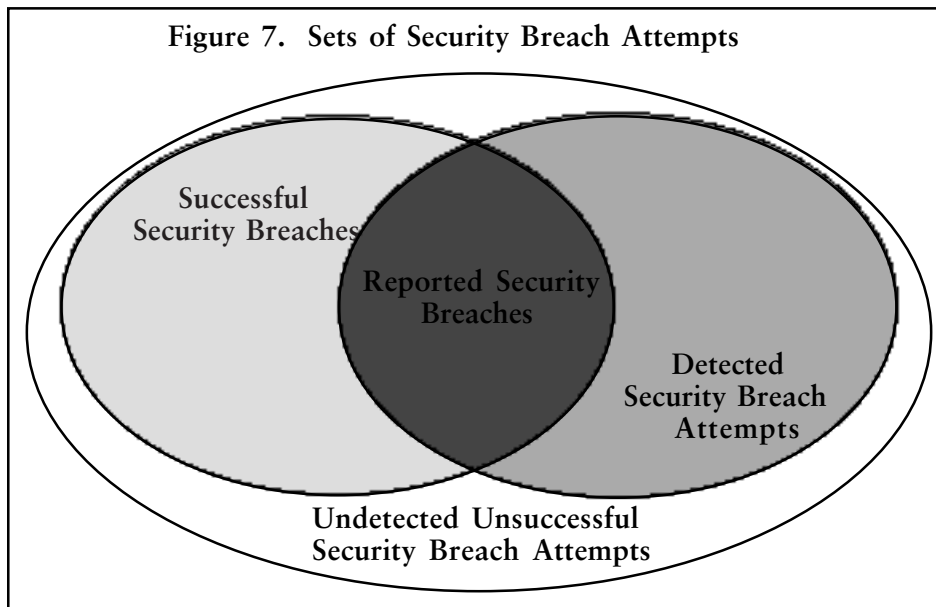
Lack of consistency in the content, conduct, and coverage of information security data collection continues to confound measurement efforts. Practical implementation issues aside, no consensus yet exists on the specific quantities to be monitored, as evidenced by the distinct lack of consonance among computer security surveys. Of the three most prominent surveys, conducted by *Information Week* magazine, *Information Security* magazine, and the Computer Security Institute, none claims to be statistically representative of any population or group. They are, at best, anecdotal evidence of computer security postures and incidents, as perceived by the survey respondents. The lack of standardization in terminology, valuation, and classifications of security breaches results in inconsistent respondent interpretations. Because the surveys neglect to define carefully their terminology and rules for accounting incidents, costs, and consequences, the reliability of the surveys as an accurate snapshot of security is suspect.

Articulating metrics for computer security risk is not a completely insurmountable task. As Figure 5 demonstrates, the ALE-based risk model identifies five key variables for which data should be obtained: frequency of bad events, consequences of bad events, measures of safeguard efficacy, costs of implementing safeguards, and additional profits enabled by safeguards.

Even if most people agreed with the selection of these variables, serious implementation hurdles remain that could undermine the credibility of data collection efforts. Previous attempts to characterize the frequency of actual and attempted computer security breaches met with fundamental uncertainty about the reliability of the gathered statistics. In Figure 7, the universe of security breach attempts is broken into four overlapping sets: Unsuccessful Undetected Security Breach Attempts, Successful Security Breaches, Reported Security Breaches, Detected Security Breach Attempts. Only the detected breaches and reported security breaches are measurable.⁵⁰

⁴⁹ President's Commission on Critical Infrastructure Protection, *Critical Foundation: Protecting America's Infrastructures* (Washington, DC: U.S. Government Printing Office, 1997), p. 27.

⁵⁰ The rationale for counting breach attempts revolves around the estimation of effectiveness for protective safeguards which prevent such attempts from turning into successful intrusions.



Thus, the relative sizes of the undetected breaches and attempts are highly uncertain. Definitional issues associated with classifying activities as either security breaches, attempts, or typical user behavior further complicate these measurements.

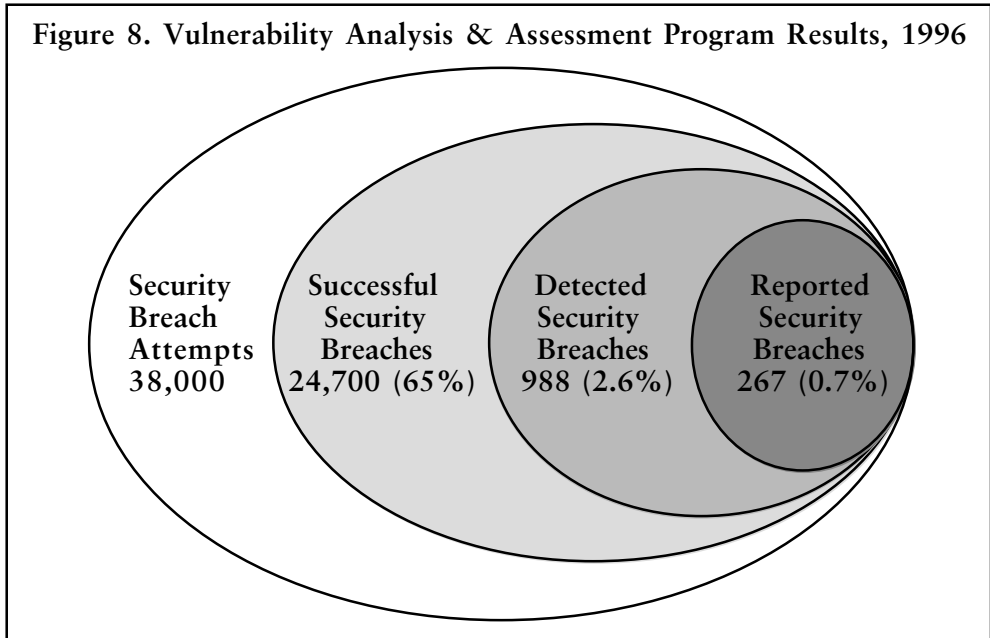
The Defense Information Systems Agency (DISA) instituted its Vulnerability Analysis and Assessment Program in an attempt to estimate the relative sizes of the sets shown in Figure 7.⁵¹ In July 1996, the agency issued its one and only publicly distributed report on this ongoing program's results. The report estimated that 96 percent of the successful break-ins were undetected, and, of the few that were detected, only 27 percent were reported; see Figure 8 for a summary of the program results as of 1996.⁵²

The implications of this report are twofold. First, they suggest that estimates of computer intrusion⁵³ activity, based on reported data, may severely underestimate the true numbers of actual and attempted computer security breaches. Second, the fact that DISA has been unwilling to release subsequent results from this continuing program demonstrates the tremendous power of suppressive forces to prevent information sharing.

⁵¹ The Vulnerability Analysis and Assessment Program is essentially a very large simulation/red-teaming exercise that utilizes publicly available "hacking" algorithms and tools to attempt to breach the computer security of Department of Defense information systems.

⁵² U.S. General Accounting Office, *Information Security: Computer Attacks at Department of Defense Pose Increasing Risks* (Washington, DC: U.S. Government Printing Office, May 1996), p. 19.

⁵³ I use the phrases "computer crime," "computer intrusion," "computer security breach," and "computer security incident" interchangeably.



3.2 Diffidence about Data Sharing

Motives for institutional reluctance to share data on computer crime activity range from concerns about legal liability and loss of competitive advantage to worries about tarnished reputation and embarrassment. Prevailing strategies today for avoiding legal liability appear to encourage companies to adopt a “hear no evil, see no evil” approach to information gathering. Fearful of being found negligent in the event of an information security breach, some companies, on the advice of legal counsel, intentionally avoid collecting computer crime data. The legal counsels argue that if the company were sued for damages resulting from a computer security breach, any collected data could be subpoenaed and used against the company to prove that it had a foreknowledge of computer risks and took inadequate precautionary measures. In the absence of data, however, the plaintiff’s task becomes significantly more difficult, especially if the company can cite industry custom as its defense.⁵⁴

This strategy, however, is no panacea. Its success hinges on the lack of comparable computer crime data. If such data were to be obtained from another, similar organization, then the plaintiffs may have a basis for a credible case, and the defendants, along with their accountants, may find themselves fighting an additional charge of poor accounting practices that failed to document and account for computer-related losses.

Companies in some competitive industries, such as banking, electronic commerce, and high technology, are especially sensitive to the prospect of losing market position as a result of major computer security incidents and customer perceptions of information security weakness. The tangible losses from diminished reputation that can be caused by public disclosure of a computer security failure have already taught businesses painful lessons about the ruthlessness of the marketplace and the importance of customer

⁵⁴ See previous discussion on liability in Section 1.6, Underlying Forces for Change.

confidence in business information security.⁵⁵ While assurances about identity protection and facilities for anonymous reporting may ease some of these concerns, businesses remain very reluctant to participate in information sharing.

Sensitivity to the public perception of computer security strength is not the exclusive province of private enterprise. Government agencies and affiliates are also extremely careful about managing the public's perception of their information security. Lawrence Livermore National Laboratories underwent extensive computer security training and awareness exercises during the summer of 1999 as a result of concerns expressed by congressional authorities about security at the national laboratories in general. The increased level of scrutiny reinforces the already entrenched propensity not to share computer intrusion data.

Government agencies such as the Department of Defense and the national laboratories are among the most popular targets for Internet-based computer hacker activity.⁵⁶ Because general security, above and beyond computer security, is a core competence and emphasis in these organizations, they are in a unique position of having both the capacity and the incentive to collect and to analyze computer security data. Their disinclination to share that data makes other government efforts to institute nationwide information-sharing seem somewhat disingenuous by indirectly reaffirming private organizations' tendencies not to share.

3.3 Relevance of Data

Some computer security experts have argued that the collection and analysis of past computer security intrusions is a pointless waste of resources. They point out that

There are no statistically valid samples applicable to specific vulnerabilities for use in information security risk assessment. Because future frequency is unknown and cannot be adequately estimated, we attempt to use data on past loss experiences. But past data consist of a series of known events; they are not a complete or valid sampling set of independent observations —as demanded by the laws of probability.⁵⁷

Valid statistical sampling aside, these experts would further argue that uncertainties in human behavior, changes in technology, growth of the Internet, and dubious asset loss calculations make past data irrelevant to future trends in computer criminal activity.

These arguments, while persuasive to some, are ultimately not convincing. They are predicated on an expectation that past data must have near perfect predictive powers. Many similar arguments could have been made against the collection of mortality statistics

⁵⁵ The Citibank incident in 1995 figures prominently as a classic example of a company announcing an information security incident to the public only to see its competitors immediately lobby its major clients to switch banks. See Saul Hansell, "A \$10 Million Lesson in the Risks of Electronic Banking (Hacker Taps into Citibank Accounts)," *New York Times*, vol. 144, August 19, 1995, pp. 15(N), 31(L), col 2. Rumored revenue losses incurred by Citibank as a result of its competitors' predations hover around \$100 million.

⁵⁶ The term "hacker" was first applied by computer programmers in the 1950s to refer to pioneering researchers who were constantly adjusting and experimenting with the new technology (Steven Levy, *Hackers: Heroes of the Computer Revolution*. New York: Dell Publishing, 1984, p. 7). These "hackers" tended to be unorthodox, yet talented, professional programmers. Though still in use today, this denotation is largely limited to small circles of computer professionals. When I use the term "hacker," I am specifically referring to the more popularly understood definition: someone who obtains unauthorized, if not illegal, access to computer systems and networks.

⁵⁷ Donn B. Parker, *Fighting Computer Crime: A New Framework for Protecting Information* (New York: John Wiley & Sons, Inc., 1998), p. 270.

in the early days of life insurance. Clearly, however, uncertainties in human behavior, advances in technology, and questions of valuation did not prevent the development of an industry that today constitutes an integral part of the modern risk-management landscape.

Using past data to predict the future has been likened to driving forward while only looking in the rearview mirror. Gentle curves in the road can be predicted and followed, but sharp turns are unforeseeable. Thus, the use of past data should be tempered by the understanding that it provides only a partial answer. Supplemental steps must also be taken to predict the sharp turns in the road and account for them accordingly.

The contrarians would contend that the road ahead for computer security risks has only sharp turns in it and that looking for gentle bends is, therefore, a waste of time. However, the little data that do exist tend to contradict this contention. The Computer Emergency Response Team Coordination Center (CERT/CC) at Carnegie Mellon University maintains one of the largest collections of publicly available computer incident response information. While new vulnerabilities, methods of attack, security remedies, and viruses are being logged almost daily, new incidents involving older, well-known attacks, sometimes dating back to 1989, continue to be reported with surprising regularity.⁵⁸ In an analysis of 1998 intrusion data for several government agencies, 45 percent of all computer intrusions were from such well-known attacks that they could have been prevented if current software patches and updates had been applied.⁵⁹ Thus, the road ahead may bend with human whim and technological advance, but it does not appear to bend too sharply too often.

3.4 Finding Data

The state of publicly available data on computer security risks is, as might be expected in light of the previous discussion, rather discouraging. While the various databases of computer vulnerabilities, viruses, and software fixes continue to improve,⁶⁰ statistically representative data on computer security breach attempts, successful attacks, consequences of attacks, valuation of those consequences, steps taken to ameliorate risk, and even basic resources devoted to information security remain woefully inadequate. The best public data can be found in three annual computer security surveys—conducted independently by the Computer Security Institute /Federal Bureau of Investigation (CSI/FBI), *Information Week*, and *Information Security*⁶¹—and in the American Society for Industrial Security's (ASIS) annual survey on intellectual property loss.⁶² These surveys clearly stipulate that they are not scientifically conducted, should not be mistaken as such, and are meant merely to be illustrative. Although the surveys drew responses from different populations, they paint a fairly consistent picture. Both magazine surveys appear to receive most of

⁵⁸ John D. Howard, *An Analysis of Security Incidents on the Internet 1989–1995*, Ph.D. thesis, Department of Engineering and Public Policy, Carnegie Mellon University, April 7, 1997, pp. 247–258.

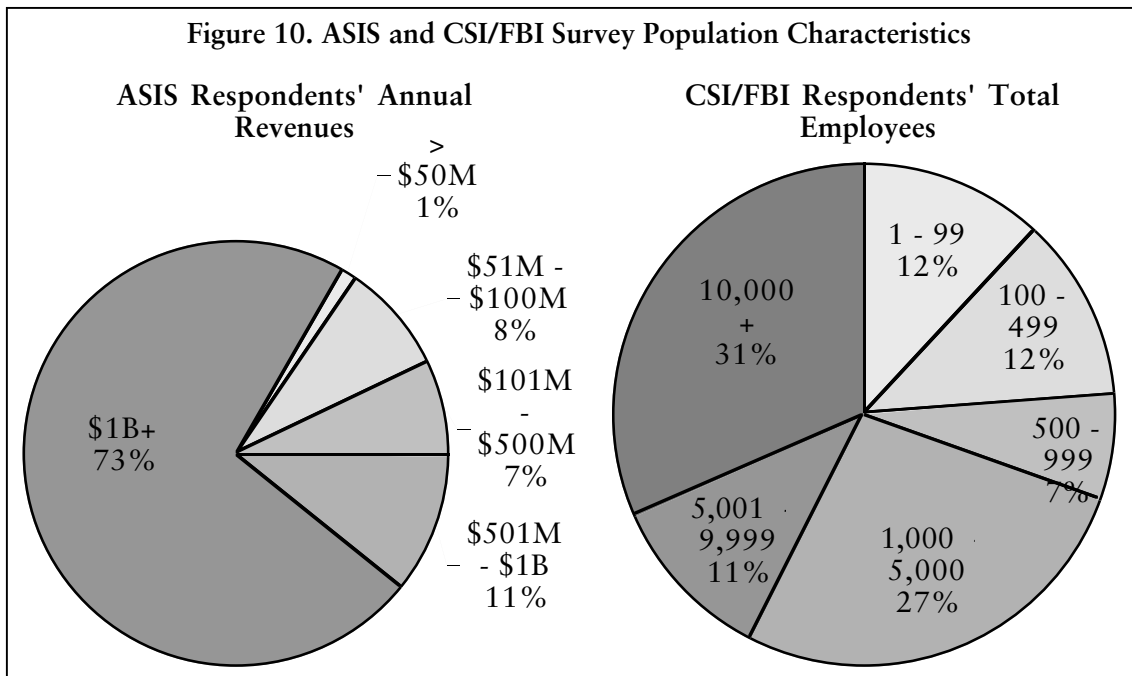

⁵⁹ As a result of persistent data-sharing diffidence, the source of this information must remain anonymous.

⁶⁰ For examples, see Computer Emergency Response Team Coordination Center at <<http://www.cert.org>>, MITRE Corporation's Common Vulnerabilities and Exposures dictionary at <<http://www.cve.mitre.org>>, Internet Security Systems' X-Force vulnerabilities database at <<http://x-force.iss.net>>, Department of Energy's Computer Incident Advisory Capability at <<http://ciac.llnl.gov>>, and Symantec Corporation's virus encyclopedia <<http://www.symantec.com/avcenter/vinfodb.html>>

⁶¹ For online survey results, see the following: <<http://www.gocsi.com/prelea990301.htm>> for the 1999 CSI/FBI Computer Crime and Security Survey, <<http://www.informationweek.com/743/security.htm>> for the *Information Week* 1999 Global Information Security Survey, and <<http://www.infosecuritymag.com/july99/cover.html>> for the *Information Security* 1999 Industry Survey.

⁶² The ASIS survey is of intellectual property theft and therefore encompasses more than theft by computer-enabled means. Its results are nonetheless instructive. For more information, see Dan T. Swartwood and Richard J. Hefferman, "ASIS Trends in Intellectual Property Loss Survey Report" (Alexandria, VA: American Society for Industrial Security, International, 1998).

their responses from smaller, private enterprises, as reflected in the information security budgets shown in Figure 9, while the ASIS and CSI/FBI surveys are biased toward larger organizations, as evident in the annual revenue and employee head-counts of Figure 10.⁶³



⁶³ As was pointed out in the beginning of this chapter, direct comparison of these surveys is often complicated by the heterogeneity in questions asked.

3.4.1 Terminology

Computer security suffers from a sometimes painfully inaccurate vocabulary. Currently, several efforts are under way to address the terminology crisis.⁶⁴ Until they are successful, definitions will likely remain imprecise. The surveys themselves do not define their terminology, leaving the respondents to assume their own definitions. Below is a list of common computer security incident terms and definitions.

Access abuse:	Inappropriate use of computer or network resources.
Active wiretapping:	Unauthorized observation of a data communications transaction.
Authorized non-employee access abuse:	Access abuse by a contractor, business partner, or other authorized person who is not an employee.
Computer virus:	A self-propagating computer program that may directly or indirectly cause one or more security incidents.
Data system integrity loss:	Tampering with an information system.
Denial of service:	Unauthorized suspension of network or computer system function(s).
Destruction of comp. resources:	Same as sabotage of data or networks.
Destruction of data:	Erasing data.
Employee access abuse:	Inappropriate use of computer system access by an employee.
Financial fraud:	The use of deceit or trickery to embezzle money.
Hacking of phone/PBX:	Same as telecom fraud (see below).
Information loss:	Removal of information possession by theft, modification, or destruction.
Insider abuse of net access:	Inappropriate use of network access by someone authorized to use the network.
Laptop theft:	Unauthorized taking of a notebook computer.
Leak of proprietary information:	Same as theft of proprietary information.
Manipulation of software apps:	Unauthorized modification of software applications.
Manipulation of system software:	Unauthorized modification of the system software.
Sabotage of data or networks:	Destroying data or suspending the function of a network.
System penetration by outsider:	Unauthorized access to a computer system obtained by someone not affiliated with the system owner.

⁶⁴ See John D. Howard and Thomas A. Longstaff, *A Common Language for Computer Security Incidents*, SAND98-8667 (Albuquerque: Sandia National Laboratories, October 1998); Donn B. Parker, *Fighting Computer Crime: A New Framework for Protecting Information* (New York: John Wiley & Sons, Inc., 1998); and Common Vulnerabilities and Exposures dictionary at <<http://www.cve.mitre.org>>.

Telecom eavesdropping:	Unauthorized observation of a telecommunication transaction.
Telecom fraud:	Theft of telecommunication services.
Theft of computer resources:	Using a computer/network for unauthorized activities.
Theft of data, trade secrets:	Same as theft of proprietary information.
Theft of proprietary information:	Unauthorized taking of sensitive or confidential data.
Trojan horse:	A hidden software module that can be activated to perform an unauthorized task; a form of virus.
Unauthorized access by insider:	Computer system or data access obtained without proper permission from the system administrator by someone affiliated with the system owner but not given permission to access the penetrated system or data.
Unauthorized access by outsider:	Computer system or data access obtained without proper permissions from the system administrator by someone not affiliated with the system owner.
Unauthorized network entry:	Network access obtained without first receiving proper permissions from the network administrator.
Virus contamination:	Computer system infection by a computer virus.


The following is a glossary of current computer security technology terms.

Access control:	Policies, procedures, and mechanisms by which users are authorized and granted privileges to use computer systems and networks.
Anti-virus software:	Computer programs that detect the presence of and remove known viruses in a computer system.
Basic user password:	Common password access control technology. Users present a user name and a password of their choosing to gain access to a computer system.
Biometric authentication:	The use of biologically distinct measures to prove user identity, for example, fingerprint scans, retinal scans, etc.
Digital certificate:	A form of digital identification in which a certificate authority vouches for the identity of the certificate presenter.
Disaster recovery:	Policies, procedures, and mechanisms by which the functionality and data of a computer system and network can be restored after a disaster.
E-mail security:	General security policies, procedures, and mechanisms to protect e-mail confidentiality, authenticity, and integrity.
Encrypted file:	Use of cryptographic techniques to scramble a file, thus denying anyone without the proper key access to its contents.


Encrypted login:	Use of encryption to scramble identification, authentication, and authorization communications.
Encryption:	Application of cryptographic techniques to scramble data so that only the proper key holder(s) may unscramble it.
External certificate authority:	Use of an outside organization to vouch for digital certificate presenters.
Firewalls:	Specialized routers that filter network traffic between two networks, effectively separating the networks.
Hardware authentication:	The use of hardware addresses or some other form of hardware identification to vouch for a user's identity.
Internal certificate authority:	Use of digital certificates where the vouching authority is a division within the same organization.
Intrusion detection systems:	Software applications that monitor computer and network activity for computer attacks, enabling them to detect incidents as they occur and sometimes prevent damage.
Multiple logons & passwords:	Permitting more than one user on a computer system, using username and password combinations for access control.
One-time passwords:	Passwords that can be used only once.
PC access-control software:	Use of software on a personal computer to control the granting of user privileges on that personal computer.
Physical security:	Policies, procedures, and physical protection mechanisms to guard information resources. For example, locks on doors, security guards in buildings, etc.
Reusable passwords:	Passwords that can be used by their owners multiple times.
Single sign-on:	Use of an access control system whereby a user is granted privileges on multiple computer systems by presenting credentials only once, such as a username and password.
Smart cards:	Electronic cards that can be presented or scanned into a system as proof of identification.
Terminal key locks:	A physical key that must be inserted into a computer terminal before the terminal will function.

3.4.2 Annual Frequencies of Security Incidents

The computation of Annual Loss Expectance (ALE) requires an estimation of the expected annual frequency of each bad event. With respect to information security, a bad event is typically the theft, destruction, modification, or denial of use of information or information systems. Thus, the estimated frequencies with which these bad events occur are often based upon the incident rates of detected computer security breaches, as seen in Figure 11. For perspective, Figure 12 gives ASIS results on the rates of intellectual property thefts and attempts at theft. The CSI/FBI data on the number of computer security incidents experienced are consistent with another estimate from NetSolve, Inc. NetSolve operates an intrusion detection system, Pro Watch Secure, for several U.S. companies. Using data gathered in 1997, they found that serious attacks occurred at rates of 0.5 to 5 times per month per customer, translating into annual rates of between 6 and 60 serious attacks.⁶⁵



⁶⁵ NetSolve, *Pro Watch Secure Network Security Survey* (Austin, TX: NetSolve Corporation, 1997).



To characterize these rates further, the three computer security surveys asked each respondent organization whether or not it had experienced one or more of several specific computer security incidents. Although the terminology differed among the surveys, similar questions have been grouped together in Table 1 along with the affirmative response rates.

Table 1. Percentage of Respondents Reporting One or More Security Incidents					
CSI/FBI Survey		Information Week		Information Security	
Theft of proprietary info	20%	Information loss	11%	Leak of proprietary info.	18%
		Theft of data, trade secrets	5%	Theft/Destruction of data	15%
Sabotage of data or networks	15%	Data system integrity loss	11%	Theft/Destruction of computer resources	23%
		Manipulation of system s/w	3%		
		Manipulation of s/w apps	6%		
System penetration by outsider	24%	Unauthorized network entry	13%	Unauthorized access by outsiders	23%
Virus contamination	70%	Computer virus	64%	Viruses	77%
		Trojan horse	8%		
Financial fraud	11%	Fraud	5%		
Denial of service	25%	Denial of service	11%		
Unauthorized access by insider	43%			Authorized non-employee access abuse	14%
Insider abuse of net access	76%			Employee access abuse	52%
Telecom fraud	13%			Hacking of phone/PBX	12%
Telecom eavesdropping	10%				
Active wiretapping	2%				
Laptop theft	54%				

In his 1995 *Computer Security Handbook*, Hutt gives a table of information security threat likelihoods of occurrence, which in turn were taken from Carroll's 1984 book, *Managing Risk: A Computer-Aided Strategy*.⁶⁶ This compilation of specific computer security incident rates dates from the late 1970s and early 1980s and is apparently the first and only attempt to collect and report such data. These data may have been appropriate for the large data centers and computer systems of twenty years ago, but they are certainly not applicable to computing today.

3.4.3 Consequences of Security Incidents

Placing a value on the damage caused by a breach of information security is a highly speculative activity. Although some of the costs associated with information assets are readily assessable, such as resources devoted to information recovery, others are not so easily quantified. For example, the value of an information asset is highly dependent upon who possesses the information. Sensitive commercial R&D information in the hands of a competitor is significantly more problematic than if it were in the hands of a Netherlands teenager. Time sensitivity can also complicate the valuation problem. For example, a password that expires in ten seconds is worthless after it has expired, but it is quite valuable during the ten seconds that it could be used to gain system access. Add to these difficulties the quantification challenges posed by intangible values, such as reputation, trust, embarrassment, etc. and the task of valuation quickly becomes a highly uncertain and potentially contentious enterprise.

Different valuation methodologies have developed in recent years.⁶⁷ These competing methodologies, however, produce dissimilar values for same asset. Some approaches look at the costs of creating/re-creating the compromised asset, others examine costs incurred as a result of a security breach, while still others try to capture all effects on both revenues and costs. Any quantification effort faces the challenge of deciding which costs and effects are appropriately attributed to an incident. For example, if a computer system were compromised and as a result its owner invested heavily to secure that system, would the security investment be a consequence of the incident? If the owner had already planned to upgrade its security anyway, does the accounting change? If an organization maintains an incident response team or an information security directorate, how should the overhead costs of those teams be divided among incidents, if at all?

The value of information analysis presented in the previous chapter may provide some guidance on this question of inclusion. Essentially, the value of information calculation compares two possible futures—one with and one without additional information—and uses some fixed objective function to measure the value difference. A similar approach could be applied to quantifying incident consequences. A comparison could be done between two possible scenarios: one in which a security incident occurs and one in which it does not occur. The differences between the two scenarios would form a basis for valuing the consequences. By providing a forward-looking, more complete picture of value, this approach recognizes and accounts for adaptive strategies and opportunity costs. For example, if a manufacturer's order-taking system suffered a one-week outage, then the value lost should not be an entire week's revenues just because the system was unable to take orders. Orders will likely still be taken and processed by some alternative, albeit

⁶⁶ See Arthur E. Hutt, Seymour Bosworth, and Douglas B. Hoyt, *Computer Security Handbook* (New York: John Wiley & Sons, Inc., 1995), pp. 3.7–3.9. See also John M. Carroll, *Managing Risk: A Computer-Aided Strategy* (Boston: Butterworth Publishers, 1984), pp.63–86.

⁶⁷ For more information on these, see Corresponding Committee for Information Valuation, *Guideline for Information Valuation*, Will Ozier, editor (Glenview, IL: Information Systems Security Association, Inc., 1993); ICAMP Report; and Richard Power, "CSI Special Report: How to Quantify Financial Losses from Infosec Breaches?" *Computer Security Alert*, October 1999, p. 1.

probably less efficient, method. The value of the consequences is, therefore, more appropriately computed by comparing the revenues that would have resulted if the system had been operational against the revenues actually realized. In fact, because information systems have proven themselves unreliable,⁶⁸ many critical applications have some back-up contingency plan or alternative mode of operation in the event of an outage. Likewise, if system administrators are compelled to spend large blocks of time dealing with security incidents and improving system security, the valuation of those incident consequences should consider what the administrators would otherwise have been doing. If the administrators were hired exclusively to improve security, then only the fraction of their time spent investigating specific incidents and repairing damage would be attributable to those incidents. Installation of additional security measures that would have been done anyway should not be counted as incident-inspired. Similar analyses can be performed for other staff members whose time must be reallocated to different tasks as a result of an information security incident.

This approach offers a simple test for deciding which costs and effects should be included in the calculus of computer security incident consequences. Costs, such as those for investigation, litigation, liability, and regulatory violations, that would never have been incurred without an incident should be included, while development costs, general maintenance costs, and capital costs should not. The valuation of system downtime or hijacked system resources could, again, be derived from a counterfactual exercise that asks what those resources would have otherwise been engaged to do to add value to the organization. Theft of research and development (R&D) information could also be viewed in a similar vein. Companies value R&D by the future expected revenue stream that the R&D will generate. Fairly elaborate decision models can be built to assess the value of an R&D project to a company.⁶⁹ Thus, in a worst-case scenario, R&D information is irrecoverably lost by one firm and obtained by a competitor. In that circumstance an opportunity has been forever lost, and the total value lost by the firm would be the R&D project's expected net present value to the company prior to being lost. All other security compromise scenarios could be considered as some fraction of this extreme. Often, the resources invested to create the compromised information are cited as an estimate of the information's value. This assignment is erroneous because the information asset's value is not linked to its past cost of creation but to its present and future contribution to the organization. Even "intangible" values such as reputation, embarrassment, etc. can be given a basis for valuation by estimating the bottom-line impact that such effects will have on customer behavior. Industries such as banking and finance that rely heavily on reputation to engender trust are more seriously affected by incidents that damage or diminish that trust than those industries that compete on price or product/service differentiation.

Because the incident loss estimates of survey respondents are not calculated using a set methodology, they are highly uncertain in both scope of coverage and accuracy of estimation. Despite this deficiency, the FBI has estimated that the total damage to U.S. businesses as a result of computer crime in 1996 was on the order of \$300 billion.⁷⁰ To put this figure in context, the ASIS reports that U.S. industry may have lost more than \$250 billion in 1997 as a result of intellectual property theft alone. Both of these estimates appear rather high when compared to the estimated total costs of all U.S. violent crime in

⁶⁸ Peter Neumann, speaking at the Conference on Cyber Crime and Terrorism, Hoover Institution, Stanford University, December 6, 1999.

⁶⁹ See David Matheson and Jim Matheson, *The Smart Organization: Creating Value through Strategic R&D* (Boston: Harvard Business School Press, 1998).

⁷⁰ National Research Council, *Trust in Cyberspace*, Fred B. Schneider, editor (Washington, DC: National Academy Press, 1999), p. 113.


1995, which was \$426 billion, and of property crime, which was \$24 billion.⁷¹ The FBI figures are also disproportionate to damage reported in the computer security surveys, with the 1999 CSI/FBI survey reporting approximately \$124 million in losses from its 163 reporting organizations and *Information Security* magazine reporting \$23.3 million in losses for the 91 organizations that could quantify their losses.⁷² A detailed breakdown of incident losses is given in Table 2 and Figure 13. The disparity between these estimates only serves to illustrate the lack of standardization in how these statistics are collected and aggregated.

Security Incident	Incidents	Lowest Loss	Average Loss	Highest Loss	Total Losses
Theft of proprietary info	23	\$ 1,000	\$ 1,847,652	\$ 25,000,000	\$ 42,496,000
Sabotage of data or networks	27	\$ 1,000	\$ 163,740	\$ 1,000,000	\$ 4,421,000
Telecom eavesdropping	10	\$ 1,000	\$ 76,500	\$ 300,000	\$ 765,000
System penetration by outsider	28	\$ 1,000	\$ 103,142	\$ 500,000	\$ 2,885,000
Insider abuse of net access	81	\$ 1,000	\$ 93,530	\$ 3,000,000	\$ 7,576,000
Financial fraud	27	\$ 10,000	\$ 1,470,592	\$ 20,000,000	\$ 39,706,000
Denial of service	28	\$ 1,000	\$ 116,250	\$ 1,000,000	\$ 3,255,000
Virus contamination	116	\$ 1,000	\$ 45,465	\$ 1,000,000	\$ 5,274,000
Unauthorized access by insider	25	\$ 1,000	\$ 142,680	\$ 1,000,000	\$ 3,567,000
Telecom fraud	29	\$ 1,000	\$ 26,655	\$ 100,000	\$ 773,000
Active wiretapping	1	\$ 20,000	\$ 20,000	\$ 20,000	\$ 20,000
Laptop theft	150	\$ 1,000	\$ 86,920	\$ 1,000,000	\$ 13,038,000

⁷¹ Perspective, "Crime's Cost," *Investor's Business Daily*, May 9, 1996.

⁷² Because neither of these organizations considers its surveys to be statistically representative of any population, much less all of U.S. industry, they do not extrapolate from their survey results an estimate of total U.S. losses.

Figure 13. Information Week Report on Security Incident Consequences



3.4.4 Safeguards

Ideally, safeguard selection would be based on a balancing of cost and effectiveness. The specific capital, installation, and maintenance costs of safeguards can be readily estimated from a host of vendors today. These vendors offer a full range of security services from basic security technology to implementation of security solutions to the complete management of information security.⁷³ Third-party comparisons of vendor pricing and services can be found in computer security journals, such as *Information Week* and *Information Security*, from professional organizations such as the Information Systems Security Association and CSI, and from proprietary research companies such as International Data Corporation, Forrester Research, and the Gartner Group.

In stark contrast, the measurement of safeguard efficacy remains primitive in a very few cases and completely elusive in the rest.⁷⁴ At present, no formal methodology exists for rating the security of one computer against the security of another. The difficulty of the measurement problem derives from two inherently uncertain quantities: prediction of the attack profile and estimation of security policy compliance. Both factors depend upon humans who have a diverse range of incentives and upon technology that is constantly changing. Attackers' motives, including thrill-seeking, twisted altruism, personal pride, revenge, and monetary gain, inspire varying degrees of persistence while information asset owners' and users' incentives may not always motivate them to follow security procedures faithfully. Technology for both facilitating and repelling attacks continues to evolve and improve at uneven rates. Both of these uncertainties severely complicate efforts to develop reliable measures of safeguard efficacy.

⁷³ Quantification of the full impact that security measures may have on an organization, especially measures that require changes in worker behavior and in their access to information, is not as easily calculable.

⁷⁴ Ratings processes do exist for some specific technologies, such as firewall, anti-virus software, and intrusion detection systems. However, these ratings are neither comparable across technologies nor applicable in an information security risk assessment context.

Accepting the uncertainties and capturing them with probability distributions is one way to bound the problem of efficacy and prevent it from derailing the entire risk-management process. One suggestion for estimating efficacy is to examine the information security practices and security incident rates of several similar organizations. Such a study might yield results if patterns could be found that indicate comparatively better or worse security practices. Although such patterns would not constitute definitive measures, they would be a credible basis for rating efficacy and choosing safeguards. However, since neither the data nor the cooperation of very many organizations is likely in the immediate future, data on safeguards remains largely confined to technology adoption rates, like those found in the survey results of Figures 14, 15, and 16.

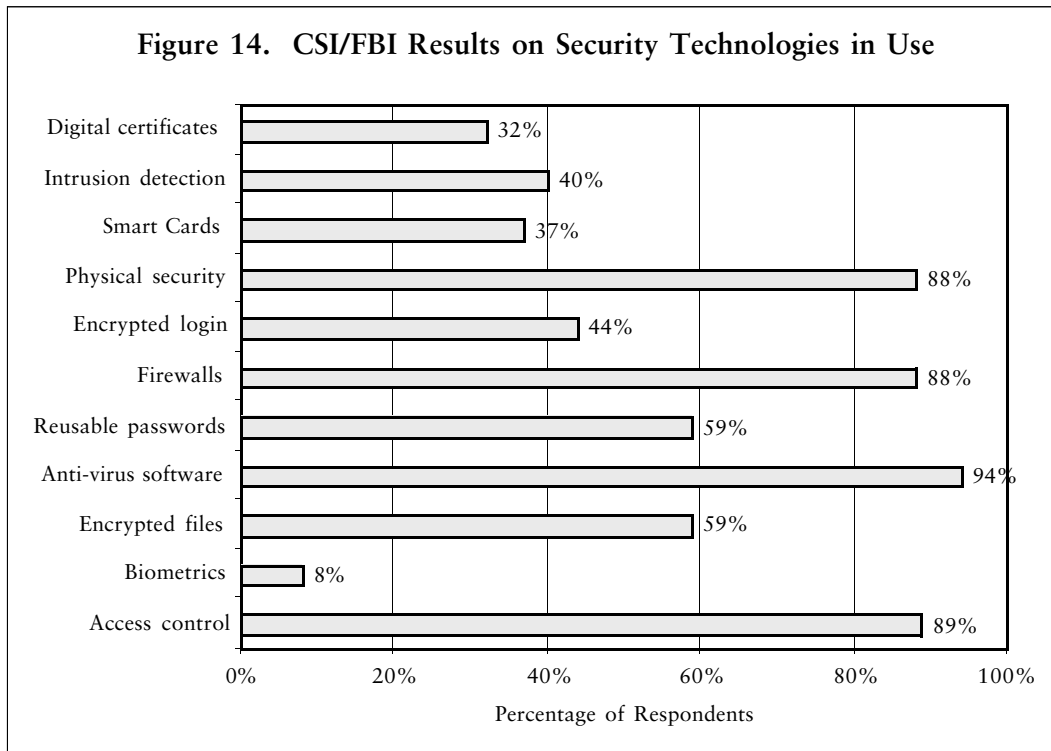


Figure 15. Information Security Report on Technologies in Use


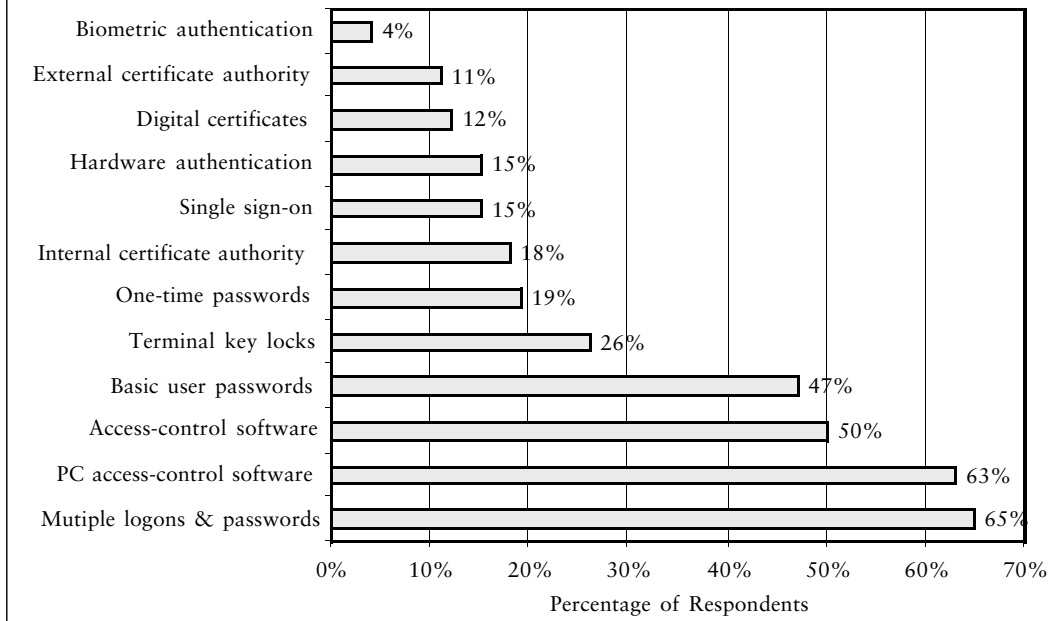


Figure 16. Information Week Results on Technologies in Use



3.5 Summary

Clearly, improvements in data collection, standardization, and dissemination are needed. The little data presented here represent the best that are publicly available for quantitative characterization of computer security. Although organizations like SRI International's I4 offer venues for confidential information sharing among its members, even these forums do not systematically gather, analyze, and disseminate any quantitative information on computer security risks. Given that organizations today do not have the luxury of waiting until this data crisis is resolved, an example is presented in the following chapter to show how uncertain data, like the data presented here, can be used in the decision analysis framework of Chapter 2 to explicate the implicit trade-offs of a computer security risk-management decision and to provide meaningful guidance on the question: How much is enough?

Chapter 4 Example Model and Analysis

Example is always more efficacious than precept.—Samuel Johnson⁷⁵

The following case study is illustrative only and does not reflect the actual security posture or decisions of any specific organization. Rather, it is a hypothetical example, reflective of the situation in which many organizations might find themselves today. The rationale behind presenting the study is twofold. First and foremost, a demonstration that marries the general approach proposed in Chapter 2 with the data presented in Chapter 3 should help the reader to understand better the benefits and limitations of both. The exercise will also show how highly uncertain data and expert judgments can be used to ground a credible analysis, highlighting the relative importance of critical variables and explicating the implicit value judgments of a computer security risk-management decision. Second, the example also illustrates strategies for resolving implementation issues that inevitably arise when modeling.

4.1 First Iteration

4.1.1 Parameters

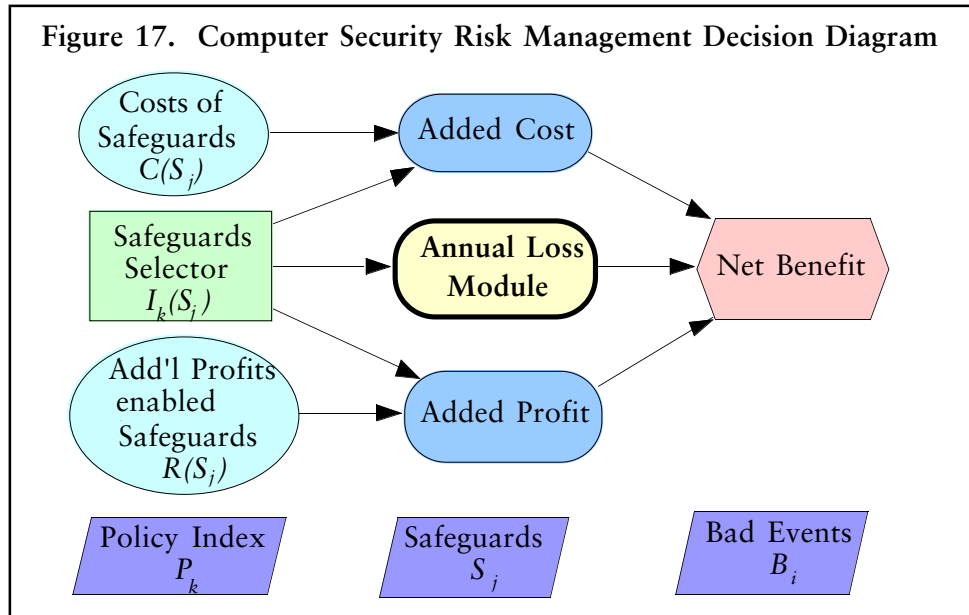
A baseline for comparative analysis must first be established. In a real-world analysis, the baseline would be the organization's status quo security posture, as described by historical data, industry reports, and internal expert assessment. For this example, we will construct a hypothetical company similar to the ones that participated in the 1999 CSI/FBI Computer Crime survey. This company shall be a large, high-technology-oriented one with some 10,000 employees and annual sales exceeding \$500 million. Most of its internal network is separated from the public network by a general firewall. Anti-virus software is not uniformly deployed throughout the organization, and virus definitions are locally updated only intermittently. Access control, including identification, authentication, and authorization, is locally controlled by individual groups, departments, and administrators. Perimeter physical security is present with facility access regulated by badge identification.

Following the diagram presented in Figure 17 as our roadmap, our next step is to bound the analysis by assigning values to the index variables. The *Policy Index* holds the labels for all security strategies, or baskets of safeguards. Initially, those labels shall be arbitrarily set as the following: Status Quo, Minor Improvement, Major Improvement, and Maximum Improvement. The *Bad Events* index will include the following general breaches of computer security:

Information Theft:	Unauthorized taking of valuable information.
Information Modification:	Unauthorized change of valuable information.
Information Destruction:	Unauthorized destruction of valuable information.
System Outage:	Removal of an information system from active service.
Employee Theft:	Unauthorized taking of money or other non-information assets by an employee for personal use.
System Degradation:	Diminished performance of an information system.

⁷⁵ Samuel Johnson, *The History of Rasselas, Prince of Abyssinia: A Tale* (London: Printed for John Sharpe, 1817).

Figure 17. Computer Security Risk Management Decision Diagram



A description of the *Safeguards* under consideration is listed below. These safeguards are somewhat general in nature, as is appropriate for the first iteration in the analysis cycle, and represent new measures that the company is considering adopting.

Security Awareness:	Instituting a large-scale security awareness program aimed at employees to encourage them to improve their security practices, such as using better passwords, locking computers when not in use, watching for suspicious behavior, etc.
HW/SW Network Upgrade:	Upgrading all software and hardware connected to the network with the latest versions and security patches.
Response Team	Creating a computer security incident response team to prevent, respond to, and recover from security breaches.
Nightly Back-ups:	Instituting automated, organization-wide, nightly back-ups of all important information stored on computer systems.
Encryption:	Widespread deployment and use of encryption technology to protect both communications and stored information assets.
Central Access Control:	Development of a centralized system for managing and tightly controlling all access control functions.
Firewalls:	Segregation of internal network using firewalls for each major division within the organization.
Screen Locking Software:	Widespread deployment of software that automatically locks a computer's screen when the system is not in use.
Security Management Team:	Creation of a high-level management team to coordinate computer security for the entire organization.

Comm Content Screening:	Comprehensive filtering of all electronic communications into and out of the organization to prevent the loss of valuable information assets.
Anti-Virus Software:	Company-wide deployment of anti-virus software that will automatically update virus definitions as they become available.
Intrusion-Detection System:	The implementation of an intrusion-detection system by hiring a network management/monitoring firm.

4.1.2 *Input Variables*

The nodes on the left sides of the diagrams in Figures 17 and 18 represent the model input variables. These variables are discussed below with initial value range estimates. Unless otherwise noted, a variable with a range of possible values will be assigned a uniform probability distribution over that range in the model.

The *Costs of Safeguards* are the annual implementation costs of each safeguard. Point estimates have been obtained for this first pass through the analysis. The rationalizations for each cost estimate are given below. Generally, the figures include worker time and direct costs of equipment, software, and services to implement the safeguard. The costs are ballpark figures culled from sales quotes, industry survey reports, and expert estimates for a large high-tech firm. Intangible costs, such as effects on employee morale, customer confidence, and opportunity costs, have been purposely ignored for this first pass through the analysis. Thus, these cost estimates are conservative ones with the “real” cost to the organization likely being higher.

Security Awareness:	\$200,000–\$600,000 The production, distribution, and overhead costs of the program are assumed to total roughly \$1-3 million, amortized over 5 years.
HW/SW Network Upgrade:	\$500,000–\$600,000 The company is assumed to have 100 servers, each requiring about 6-12 hours of time from an experienced system admin, who is paid a fully loaded salary of roughly \$100 per hour. The rest of the computers are likely to need minimal attention, on the order of half an hour.
Response Team	\$600,000–\$1,000,000 The response team is assumed to consist of 3-5 people being paid an average, fully loaded annual salary of \$200,000.
Nightly Back-ups:	\$560,000 Backup software costs roughly \$400 per server. The organization is assumed to have 100 servers and the process to require 1 hour a week of system admin time per server.
Encryption:	\$1M–\$1.1M The software costs of encrypting e-mail communications and hard drive data range from \$50 to \$60 per user plus .5 hours of system admin time for installation.

Central Access Control:	\$8–12 Million The total cost of such an undertaking, including the costs of software development, hardware, user training, implementation, and maintenance, is roughly \$40–\$60 million. This total has been spread over 5 years, the expected lifetime for such a system.
Firewalls:	\$288,000–\$420,000 Installation and annual maintenance of 10 firewalls to protect various portions of the internal network cost about \$2400–\$3500 per month per firewall.
Screen Locking Software:	\$10,000 Nearly all personal computer operating systems already offer this feature. Thus, the cost of activating this feature is fairly trivial, \$1 per machine.
Security Management Team:	\$900,000–\$1.3 Million The management team is assumed to consist of a director (\$400,000 fully loaded annual salary) and a small staff of 2-4 people (\$200,000 fully loaded annual salaries on average) with an annual budget of \$100,000 for security activities.
Comm Content Screening:	\$65,000–\$70,000 The software and hardware costs associated with a communications screening system, such as those employed to screen e-mail, are about \$60,000 a year plus system admin time of 1-2 hours a week.
Anti-Virus Software:	\$150,000 Norton Anti-Virus software costs \$20 per user plus \$10 per user for upgrade insurance. Assume also that company-wide, uniform distribution and installation of upgrades takes place via the 100 servers and requires .5 to 1 hour/month/server of admin. time.
Intrusion Detection System:	\$200,000–\$300,000 Management/monitoring costs for a 100–200 site network by a third-party vendor run \$200,000 to \$300,000.

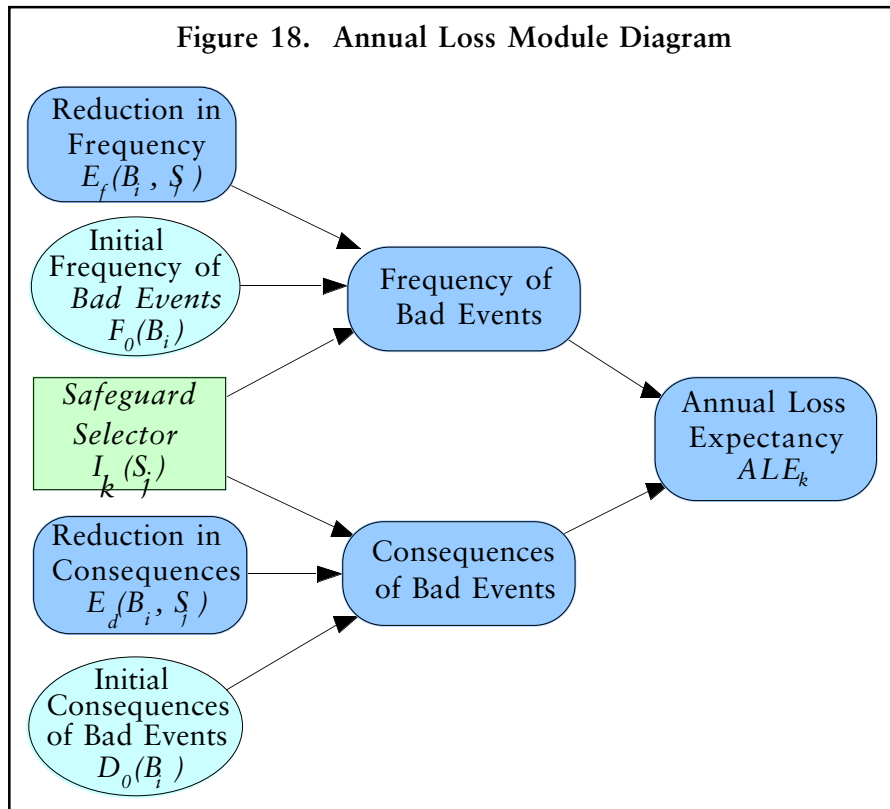
The *Add'l Profits Enabled by Safeguards* will be kept at zero for now. The results of an entire new-business financial analysis should rightly be placed here, but such an exercise would distract from the primary mission of weighing computer security risks. The potential for additional profits is certainly an important factor, and to the extent that the costs of safeguards might exceed their benefits, additional profits from new business ventures could offset those costs and become material to the decision. The advent of electronic commerce with its promise of lucrative business expansion has made this factor a particularly compelling argument for many companies to improve their security.

The *Safeguard Selector* is a switch that allows the user to include different safeguard combinations in each policy alternative. As shown in Table 3, the 1s, indicating inclusion, and 0s, indicating exclusion, under each column heading show the safeguard composition

of each policy option. Because the safeguard costs and benefits are estimated relative to the status quo, the status quo column has only 0s in it.⁷⁶

Safeguard	Status Quo	Minor Improvement	Major Improvement	Maximum Improvement
Security Awareness	0	0	1	1
HW/SW Network Upgrade	0	1	1	1
Response Team	0	0	0	1
Nightly Back-ups	0	1	1	1
Encryption	0	0	1	1
Central Access Control	0	0	0	1
Firewalls	0	0	1	1
Screen Locking Software	0	1	1	1
Security Management Team	0	0	1	1
Comm Content Screening	0	0	0	1
Anti-Virus Software	0	1	1	1
Intrusion Detection System	0	0	0	1

The diagram for the *Annual Loss Module*, as seen in Figure 18, reveals four more input variables that are needed to calculate the annual loss expectancy.



⁷⁶ Optimal safeguards selection will be discussed in the analysis section below.

An estimate of safeguard efficacy is essential to any cost-benefit calculation. In Figure 18, safeguard efficacy is represented with percentage reductions in frequencies and consequences of bad events. If available, past incident data could be analyzed to derive these reduction factors. Otherwise, expert judgments would be used to estimate the efficacy rates. Tables 4 and 5 contain matrices of such expert opinions. Table 4 shows the fractional reductions in bad event frequencies, while Table 5 shows the fractional reductions in bad event consequences. Although given as point estimates, these values will be varied from one-half and two times the listed value in model sensitivity analysis to determine whether such differences materially affect the decision.

	Info Theft	Info Mod.	Info Destr.	System Outage	Employee Theft	System Degrad.
Security Awareness	0.35	0.3	0.3	0.05	0.6	0.5
HW/SW Network Upgrade	0.45	0.45	0.45	0.45	0	0.45
Response Team	0.4	0.4	0.4	0.4	0	0.2
Nightly Back-ups	0	0	0	0	0	0
Encryption	0	0	0	0	0	0
Central Access Control	0.3	0.15	0.15	0	0.5	0
Firewalls	0.75	0.75	0.75	0.75	0.2	0.1
Screen Locking Software	0.15	0.2	0.2	0	0.4	0
Security Management Team	0.5	0.5	0.5	0.5	0.5	0.5
Comm Content Screening	0.75	0	0	0	0.3	0
Anti-Virus Software	0	0.35	0.4	0	0	0.4
Intrusion Detection System	0.51	0.51	0.51	0.51	0.25	0.51

	Info Theft	Info Mod.	Info Destr.	System Outage	Employee Theft	System Degrad.
Security Awareness	0	0	0	0	0	0
HW/SW Network Upgrade	0	0	0	0	0	0
Response Team	0	0.2	0.2	0.7	0	0.65
Nightly Back-ups	0	0.6	0.95	0	0	0
Encryption	0.95	0.95	0	0	0	0
Central Access Control	0	0	0	0	0	0
Firewalls	0	0	0	0	0	0
Screen Locking Software	0	0	0	0	0	0
Security Management Team	0	0	0	0	0	0
Comm Content Screening	0	0	0	0	0	0
Anti-Virus Software	0	0	0	0	0	0
Intrusion Detection System	0	0	0	0	0	0

Because these reduction estimates depend upon the organization and its security posture, they may not be applicable to other security assessments. The advantage of this approach, however, is that organization-specific security characteristics may be factored into the model at this stage. For example, some organizations might be insulated by a sort of “security through obscurity” because only a handful of people are familiar with their proprietary operating system, while others might enjoy general anonymity because they are small, low-profile organizations. These factors affect both the ability of safeguards to improve security and the initial estimates of bad event frequencies and consequences.

The final two key input quantities require predictions about the future frequency and consequences of bad events. For the example, values have been taken from the CSI/FBI survey, using 1997, 1998, and 1999 incident data for Theft of Proprietary Information, System Penetration by Insiders, Sabotage, Denial of Service, Financial Fraud, and Virus. These categories roughly correlate with Information Theft, Information Modification, Information Destruction, System Outage, Employee Theft, and System Degradation. The survey gives three years of frequency data for all but the system outage category. Because the average incident rates are not consistently increasing from 1997 through 1999, triangular probability distributions, and in the case of system outage a uniform probability distribution, have been used to express uncertainty about the specific rates, as shown in Table 6.⁷⁷

Bad Events	Annual Frequency Estimate
Information Theft	Triangular (0.18, 0.21, 0.25)
Information Modification	Triangular (0.40, 0.44, 0.55)
Information Destruction	Triangular (0.13, 0.10, 0.14)
System Outage	Uniform (0.25, 0.32)
Employee Theft	Triangular (0.12, 0.14, 0.15)
System Degradation	Triangular (0.83, 0.84, 0.90)

The CSI/FBI Survey gives three 1999 loss values for each category: lowest reported, mean reported, maximum reported. These three numbers are used to generate a triangular probability distribution that describes the “average” loss per incident, as seen in Table 7.



Bad Events	Annual Frequency Estimate
Information Theft	Triangular (1K, 1.848M, 25M)
Information Modification	Triangular (1K, 103.1K, 500K)
Information Destruction	Triangular (1K, 163.7K, 1M)
System Outage	Triangular (1K, 116.2K, 1M)
Employee Theft	Triangular (10K, 1.471M, 20M)
System Degradation	Triangular (1K, 45.46K, 1M)

4.1.3 Initial Results

A preliminary analysis of the three policy alternatives specified in the *Safeguard Selector* shows that of the three, the Minimal Improvement policy, with an expected net benefit of \$558.7K, is the best one to pursue. See Figure 19 for the expected net benefit values of each policy.

The cumulative probability distribution graphs in Figure 20 demonstrate the first-degree stochastic dominance of both the Major and Minimal Improvement policies over the Maximum Improvement policy. Since the Major and Minimal Improvement policies’ graphs cross near the 97th percentile, first-degree stochastic dominance is not present. However, second-degree dominance is readily apparent.

⁷⁷ For more information on continuous probability distribution like the uniform and triangular distributions, see M. Granger Morgan and Max Henrion, *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*, 1st edition (New York: Cambridge University Press, 1990), p. 96.




4.1.4 Analysis

Model analysis can often be more insightful than the actual model results. Analysis affords us an opportunity to better understand the dynamics of the model, how uncertainty affects the safeguard selection decision, and the resultant net benefit. This process not only gives guidance for further modeling and information-gathering efforts but also for the maximum value that such activities would likely add to the decision process.

For this example, one of the first questions that immediately arises is, What is the optimal combination of safeguards? The baskets chosen in the *Safeguard Selector* were done so somewhat arbitrarily. Unfortunately, the general problem of choosing the optimal basket of safeguards is a non-linear, binary integer, unconstrained maximization problem. Although advances in the field of optimization may yet solve this problem more elegantly, the best solution at present involves a judicious application of the exhaustive search algorithm. Luckily, the tractability of exhaustive search only depends upon the number of safeguards under consideration. In the example, twelve safeguards are considered, making the total possible safeguard combinations 2^{12} or 4,096. Although sizable, this number of combinations is not prohibitively large, and an exhaustive search should be well within the capability of current desktop computers.

Some steps, however, may be taken to trim the number of combinations involved. First, prohibitively expensive safeguards whose costs exceed their expected annual savings can be safely ignored. These safeguards would never be selected in an optimal portfolio because of their unfavorable savings-to-cost ratio. In Figure 21, the savings-to-cost ratio for each of the safeguards is given. The savings were calculated by assuming that each safeguard was implemented in isolation. Note that Nightly Back-ups and Central Access Control have ratios of less than one and are therefore in the pure security-oriented focus of



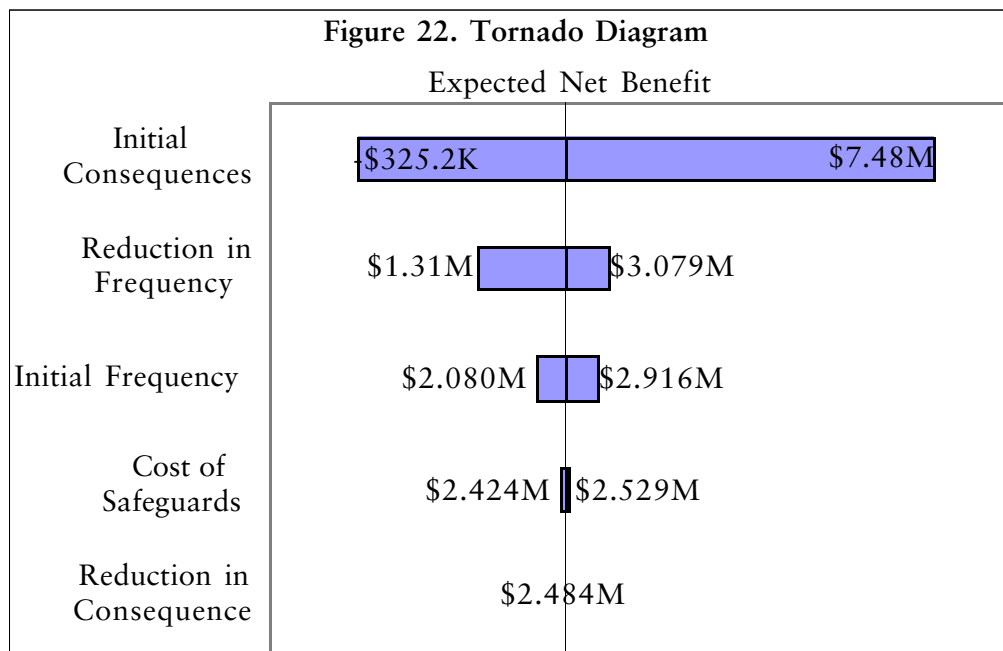
the analysis, not cost-justified. Of course, the company may have other, very good reasons to implement a nightly back-up scheme relating to protection of intellectual property from computer malfunctions or user errors.

The flip side of exclusion is inclusion, and this strategy may also reduce the number of considered combinations. When a safeguard, such as screen locking, enjoys a vastly superior savings-to-cost ratio or when the decision maker mandates the inclusion of a particular safeguard, then all combinations that exclude that safeguard can be safely excised from the optimization search algorithm. Thus, the basket of safeguards that maximizes the net benefit may be found by an exhaustive search of the remaining space. The optimal strategy for the example has a net benefit of \$2.484 million, an expected cost

of \$327,500, and includes screen locking software, communications content screening, and an intrusion detection system.⁷⁸

That most security measures were not included in this optimal solution is attributable to three factors. First, the magnitude of reported computer security-related losses does not warrant additional expenditures on more safeguards. Second, the relative effectiveness of the unselected safeguards at preventing bad events does not justify their higher costs. Third, the profile of computer security incidents and consequences is heavily weighted toward employee theft, intellectual property theft, and system degradation, and several of the unselected safeguards simply offer little protection against these bad events. Thus, if any of these underlying input values were to change substantially, the model analysis would likely return very different results. How much of a change constitutes a “substantial” amount? Both nominal range sensitivity analysis and parametric sensitivity analysis are designed to address this very question.

In nominal range sensitivity analysis, five key input variables are examined for their effect on the expected net benefit of the optimal policy: *Cost of Safeguards*, *Reduction in Frequency*, *Initial Frequency of Bad Events*, *Reduction in Consequences*, and *Initial*



Consequences of Bad Events. In instances where ranges of values are not stipulated, as is the case for both reduction factor variables, the nominal values are divided in half to obtain a low bound and doubled to obtain a high bound.⁷⁹ The analysis reveals that *Initial Consequences* has the most dramatic effect, capable of causing the expected net benefit to be as low as -\$325,000 and as high as +\$7.48 million, as shown in Figure 22.

Equally insightful is the result that both *Cost of Safeguards* and *Reduction in Consequence* have a negligible effect. This conclusion can be explained by the very low safeguard cost of the optimal policy and by the consequence reduction factors being zero for all optimal policy safeguards.

Parametric sensitivity analysis questions how much an uncertain variable must change from its nominal value to materially affect the decision. This cross-over analysis is useful


⁷⁸ No reduction of the search space was done for any of the example optimization calculations.

⁷⁹ If doubling the nominal value resulted in a high value greater than 1, it was replaced with 1.

for determining the confidence with which one may choose the optimal decision. An interesting situation exists with respect to *Initial Consequences* and *Initial Frequency* in that they share identical cross-over points. This apparent coincidence is actually evident in Equation 5 for Annual Loss Expectancy.


$$ALE_k = \sum_{i=1}^n \left\{ F_0(B_i) D_0(B_i) \prod_{j=1}^m \left[(1 - E_f(B_i, S_j) I_k(S_j)) (1 - E_d(B_i, S_j) I_k(S_j)) \right] \right\} \quad (5)$$

Since $[F_0(B_i)D_0(B_i)]$ changes uniformly regardless of whether Frequency (F_0) or Consequences (D_0) changes on a percentage basis, the cross-over points should be the same for both variables. Figure 23 illustrates the effect of fractional swings in either input variable on expected net benefit.




As shown in Figure 23, the optimal strategy of implementing screen locking, communications content screening, and intrusion detection system safeguards is the best policy to pursue provided either *Initial Consequences* or *Initial Frequency* do not drop below 44 percent of the best-guess, or mid, value or do not rise more than 37 percent above the mid value. Otherwise, the optimal low policy, which calls for screen locking and communications content screening only, or the optimal high policy, which includes screen locking, communications content screening, an intrusion-detection system, and a security awareness program, would become more attractive choices, respectively.

Similar analyses may be performed on the other two critical inputs. In the case of *Reduction in Frequency*, the optimal low policy contains the firewalls safeguard in addition to screen locking, communications content screening, and an intrusion-detection system. The optimal high policy contains only screen locking and an intrusion detection system. As can be seen in Figure 24, the cross-over points from mid to low and from mid to high are 90 percent and 195 percent of the mid value, respectively.




The optimal low policy for *Cost of Safeguards* includes all of the optimal mid safeguards plus the security awareness program. The optimal high policy includes screen locking, communications content screening, and anti-virus software. Figure 25 shows the cross-over points for *Safeguard Costs* to be roughly 95 percent and 175 percent of the mid-value.



Because many of the parameter fluctuations used in the cross-over analysis are well within the bounds of uncertainty surrounding the input variables, narrowing that

uncertainty could lead to a better decision. However, by looking at the previous three figures, one can see that, in areas where the optimal mid policy was not the best choice, the differences between the expected net benefit of the better policy and of the optimal mid were very small. Thus, the optimal mid policy is a very robust decision for the range of parametric values considered. Still, if information could be obtained that would narrow the uncertainty of one or more of the input variables, what would such information be worth? The expected value of perfect information analysis provides a useful metric for assigning the maximum value one should be willing to pay for such information. As shown in Figure 26, the *Initial Consequences* input has the highest EVOI at \$39,220. *Initial Frequency* has been purposely omitted from this analysis because the cross-over points occur outside of its established bounds of uncertainty; thus the optimal mid policy is best for all possible values.



Compared to the consequences of security breaches and to the costs of safeguards, both of which are on the order of hundreds of thousands or millions of dollars, these expected values of perfect information appear rather low. This result was foreshadowed by the parametric sensitivity analysis, which revealed that the optimal mid is a very robust decision. In essence, the CSI/FBI survey data do not reveal computer security incidents of significant enough severity and frequency to warrant the other, more expensive safeguards. The safeguards selected in the optimal policy are among the cheapest, and even when the efficacy measures of the three optimal policy safeguards are reduced by an order of magnitude, the optimal policy merely exchanges intrusion detection system and communications content screening with security awareness program and firewalls. Thus, the analysis yields the following conclusions:

1. The optimal policy will cost \$327,500 per year, is expected to save the organization \$2.811 million annually in avoided security incidents, and includes the following safeguards: screen locking software, communications content screening, and intrusion detection system.

2. The value of perfect information is greatest for *Initial Consequences* and the most that should be spent in reducing this uncertainty is \$39,220.
3. The credibility of the model results is directly dependent upon the degree of belief in the order-of-magnitude accuracy of *Initial Consequences* and *Cost of Safeguards*. It is also somewhat dependent, although to a lesser degree, on the relative efficacy estimates of different safeguards.
4. Since the *Cost of Safeguards* were lower bound estimates, a more comprehensive accounting of costs would likely exacerbate the unattractiveness of many of the already too-expensive safeguards.

The heretofore ignored factor that could play a significant role is the *Additional Profits Enabled by Safeguards*. Many businesses today are contemplating e-commerce ventures that could potentially result in significant increases in both the frequency and consequences of computer security breaches. Likewise, the ventures might also afford lucrative opportunities of profit. From a computer security standpoint, the two effects could be offsetting, with the additional expected profits more than compensating for the added security costs and increased risk exposure. However, as with all new business ventures, many factors unrelated to security make the outcome highly uncertain. With the addition of security concerns into that calculation, a planner may be forced to reevaluate and to modify the new venture strategy. In either case, the security implications of any new business venture should be considered in both the new-business model and the computer security risk management model.⁸⁰

This section has demonstrated with a hypothetical case study how an analytical model can be used to yield insights into a computer security risk-management decision. The next iteration of modeling, as indicated in the analysis calculations, should focus on the development of a better understanding of the *Initial Consequences*, *Reduction in Frequency*, and *Costs of Safeguards*. These efforts, however, should be governed by the insight that additional modeling and information is of limited value to the decision, on the order of tens of thousands of dollars at most. If the decision maker or some other important stakeholder finds fault in the order of magnitude of these variables, then adjustments should be made to their values and the model analysis should be redone before proceeding on to the next iteration of modeling and progressive refinement.

4.2 Subsequent Iterations

Ostensibly, the intent behind adding greater detail to a particular aspect of a model is to reduce the uncertainty surrounding it. Thus, the availability of information to support that modeling effort is a significant limiting factor on the aspect's extensibility. Because the foregoing analysis demonstrated fairly convincingly that further modeling efforts were of limited utility, no further example results will be presented here. Rather, a series of influence diagrams and explanations will follow to illustrate how one might extend the modeling detail behind each key variable in the model to utilize potentially available data or more accurate estimates.

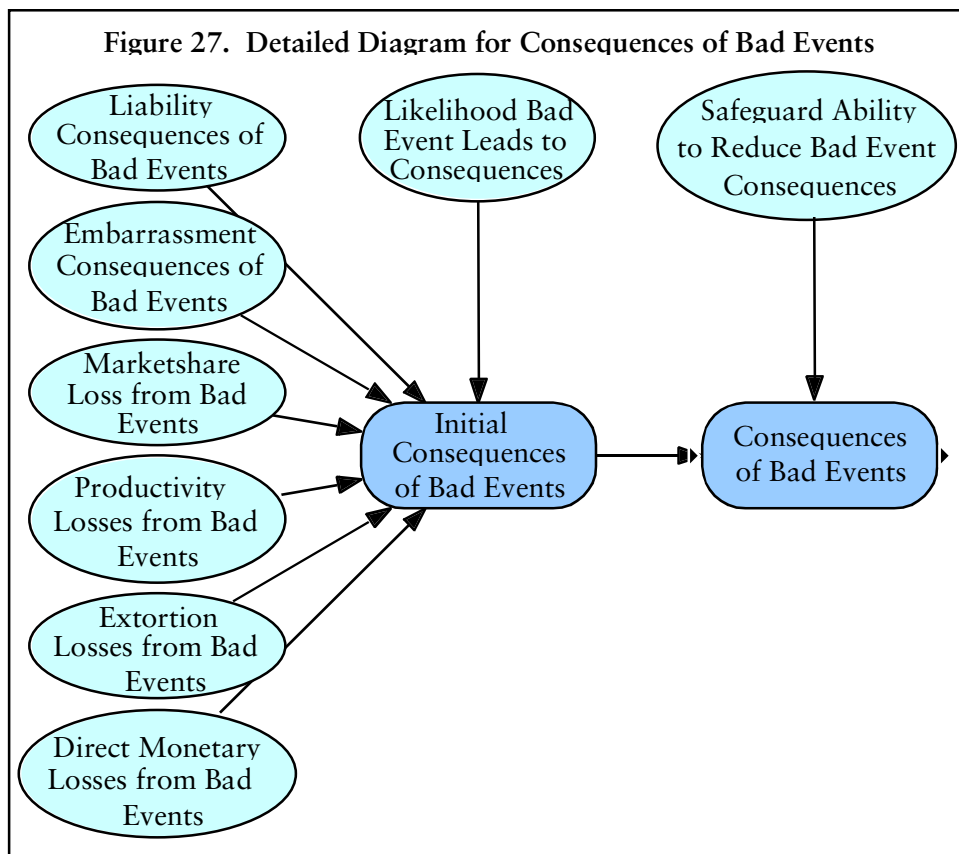
4.2.1 Consequences of Bad Events

The tasks of understanding and quantifying the consequences of a computer security breach are among the more difficult theoretical as well as practical challenges. The

⁸⁰ The rationale for leaving this factor undeveloped is somewhat self-evident. The complexity and nuances of new-business modeling are well beyond the scope of this dissertation and would represent a significant distraction from the primary purpose of computer security risk management.

difficulty rests in the counterfactual exercise required to assess the impact of a computer security breach. For example, when calculating the market-share loss resulting from a theft of trade secrets, assumptions about market development, company performance if the theft had not occurred, competitor actions, and other factors must be made to establish a baseline of comparison. The calculation also requires an estimation of who the thieves are and how likely they are to realize the value of the stolen information. As noted earlier, theft of trade secrets by a teenager in the Netherlands who is completely oblivious to their competitive importance will have a significantly lesser impact than a trade-secret theft by a rival firm. For another example, consider losses in worker productivity that might result from a computer system outage or degradation. Quantification of productivity losses requires an assessment of how well employees adapt to degraded or absent computer resources and whether productivity actually declines in an environment of limited resources. For example, the denial of Internet access for a few hours could potentially enhance the productivity of some workers who might otherwise be engaged in personal web-surfing instead focused on business matters.

The temptation to build more model detail than might be warranted is especially acute when the underlying assumptions are readily articulated. Each of the six consequence categories listed down the left side of Figure 27 could be easily developed into full diagrams of assumptions, data, and expert judgments. The progressive refinement process, however, requires small steps, and, thus, further development is probably best saved for the next iteration in the modeling process, after analysis of the current iteration warrants it.



Below is a brief description of the loss categories in Figure 27.

Liability consequences:	If lax computer security in one organization results in damages to others, that organization may be subject to liability lawsuits and forced to pay damages.
Embarrassment consequences:	Public perception of computer security strength can materially affect the prosperity and success of an organization. To the extent that computer security incidents are publicized, they might cause embarrassment and damage reputation.
Market-share loss:	If a computer security incident results in a loss of intellectual property or a delay in product development or deployment, market share could be lost to competitors.
Productivity losses:	Computer security incidents may reduce employee morale or directly hinder their ability to work, resulting in lower productivity.
Extortion losses:	Because computer security losses could be significant, the possibility exists for malefactors to attempt extortion, threatening harm to an organization unless certain conditions are met.
Direct monetary losses:	Computer-enabled embezzlement could result in direct monetary losses by an organization.

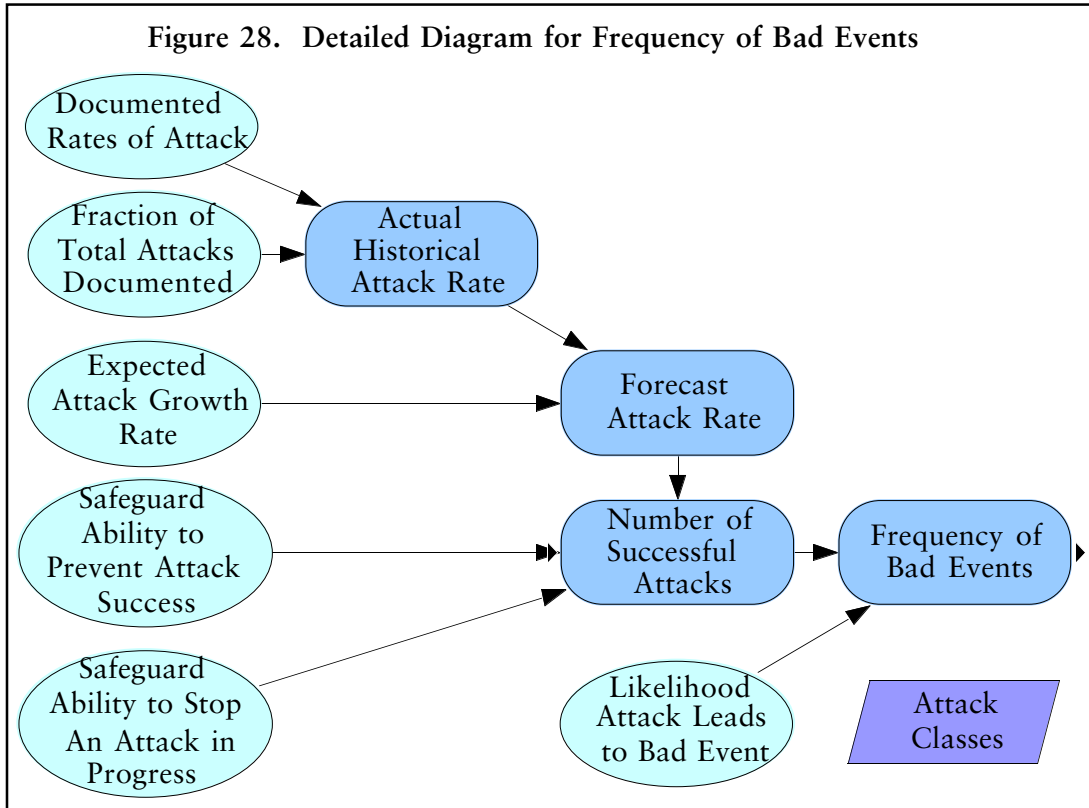
4.2.2 *Frequency of Bad Events*

Frequency of bad events data are often framed in an attack/vulnerability dichotomy. The risk model could be extended to embrace this data by creating classes of attacks. As seen in Figure 28, three new concepts are introduced into the calculation of the initial bad event frequency: attack classes, historical data, and a linkage between the new attack classes and the original bad events. Attack classes might include network-based attacks, viral infection via e-mail, social engineering, etc. The use of historical data and the explicit assumption about their reliability and accuracy could be very useful in facilitating discussions about attack frequencies in the future. The calculation of safeguard efficacy measures would, of course, need to accommodate the new attack classes distinction. Rather than efficacy against preventing bad events, the measures now look to prevent attacks from occurring, as would be case when security holes are patched, or arrest those attacks while in progress, as with intrusion-detection systems.

Figure 28 shows how these new distinctions of attack classes and historical data can be folded into the decision diagram for the frequency of bad events. Below is a brief description of the new uncertainties.

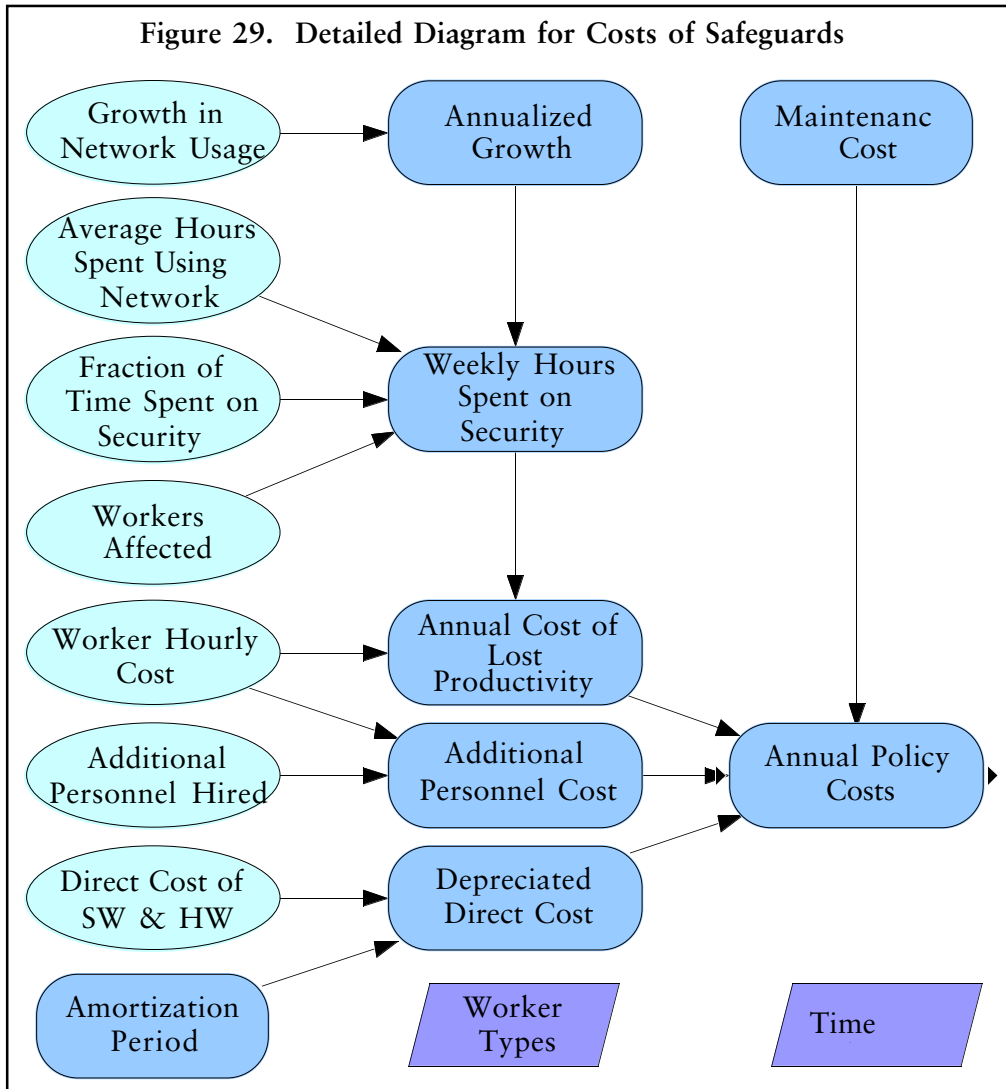
Documented rates of attack:	Historical data on the annual frequencies of different attacks.
Fraction of total attacks documented:	Estimate of the attacks reflected in the historical data as a fraction of the actual number that took place.
Expected attack growth rate:	Expected annual growth in number of attacks.

- Safeguard ability to prevent attack success: Efficacy measures of safeguards' ability to prevent an attack from being successful.
- Safeguard ability to stop attack in progress: Efficacy measures of safeguards' ability to stop an attack in progress.
- Likelihood attack leads to bad event: Likelihood that each type of attack will lead to each category of bad event.



4.2.3 Costs of Safeguards

The total cost to an organization of safeguards is potentially much greater than the direct costs of hardware, software, and system administrator time examined in the initial analysis. As was the case with valuing bad event consequences, assumptions could be made about employee dependence upon information resources and how that dependence affects morale and productivity. The implementation of some safeguards could adversely affect worker productivity by slowing computer system or network performance, mandating new security procedures, or disenfranchising employees who, out of ignorance, might resent the extra effort required. Figure 29 captures these concepts in a simplified framework and articulates the basic cost categories that made up the initial cost estimates: hardware and software costs, additional personnel hired, and maintenance costs.



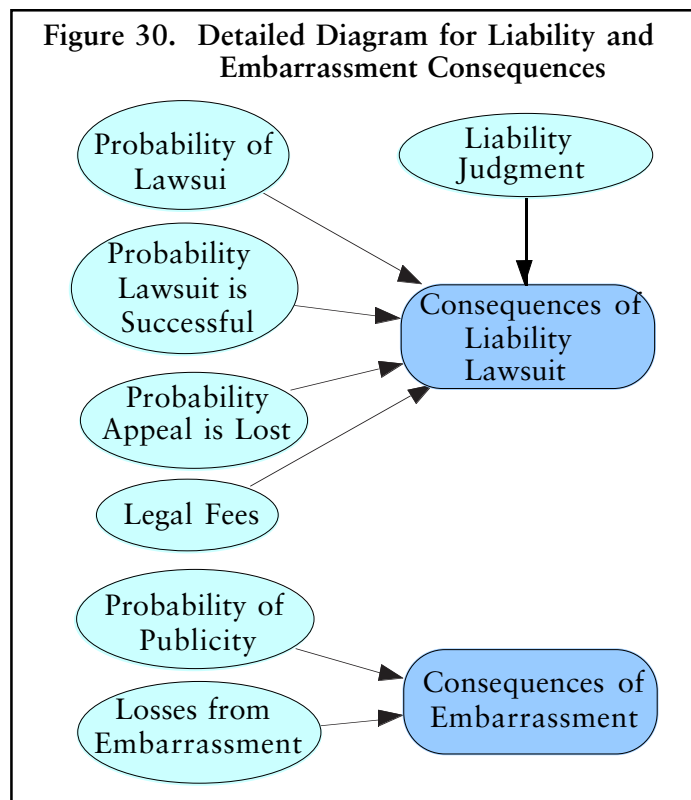
4.3 Model Adaptability

Although the example presented in this chapter was for a generic, large, high-technology company, the basic methodology and analytical model are adaptable to virtually any private organization or public agency. Conceptually, the key variables remain the same, but their dimensions and values will almost certainly change. For example, a financial institution might be more concerned about embezzlement or computer-enabled theft, whereas a defense contractor might worry most about its security image in the eyes of its primary customer, the U.S. Department of Defense. Financial institutions, consumer goods manufacturers, and others for whom reputation and public image can dramatically affect their viability would likely consider the public's perception of their security to be of great importance and consequently would want to model it in a cost-benefit analysis of security. Techniques are presented in this section for adapting the modeling framework to account for these different concerns and values.

To the extent that values and consequences can be expressed in monetary terms, an analysis may proceed as in the preceding example. Since the implications of computer security violations can often be reduced to bottom-line financial impacts, dollars are a

convenient metric for measuring costs and benefits. Even “intangible” values, such as reputation, and uncertain damages, such as liability judgments, can be quantified by using probabilistic estimates of their effects upon an organization’s present and future prosperity.⁸¹ For example, if a serious security incident at a bank were to become public knowledge, the bank might estimate that as many as one-quarter to one-third of its major customers would take their accounts elsewhere. Then, after a period of confidence-building, the bank might guess that it could win back some or all of those customers. Thus, the impact of the security incident and the subsequent bad publicity could be captured in a discounted-cash-flow, real-options, or other quantitative business model.⁸² The output of that model would then be used as an input to a risk assessment.

At times, the consequences of a computer security breach may not automatically result from the event itself. One or more intermediate events may also be necessary for the full brunt of the consequences to be felt. For example, damages from liability lawsuits are only realized after a suit has been brought, the trial lost, and the appeal denied. Likewise, damages from public embarrassment or loss of reputation are contingent upon the media discovering the embarrassing event and publicizing it. Figure 30 shows one way to model the uncertainty of whether or not certain consequences will come to pass after computer security has been breached. The *Consequences Liability Lawsuit* node represents the probabilistic expectation of losses from a liability lawsuit, while the *Consequences Embarrassment* node represents losses from bad publicity and public embarrassment.



⁸¹ Net-present-value financial analysis is one useful technique for estimating the value of future revenues and losses to a business; for more information, see Robert G. Cooper, Scott J. Edgett, and Elko J. Kleinschmidt, *Portfolio Management for New Products* (Reading, MA: Addison-Wesley, 1998).

⁸² For more information on business modeling, see Lenos Trigeorgis, *Real Options: Managerial Flexibility and Strategy in Resource Allocation*, 4th printing, 1999 (Cambridge: MIT Press, 1996).

Some organizations, however, may hold values that are not readily converted into monetary terms. The U.S. Department of Defense (DOD), for example, might be very concerned about military readiness and saving lives, two values to which many people are hesitant to attach dollar equivalents. A cost-benefit analysis, however, requires that some trade-off be made between the costs of security and its benefits. Thus, in the case of the DoD valuing readiness, a metric for readiness would be needed to gauge both the benefit of security measures and the costs of security breaches. Provided the metric is quantitative, a variant of the proposed modeling framework could be used in a dual-track assessment of safeguard costs and benefits. One track of the analysis would weigh units of readiness, while the other would be in regular dollars. Thus, all assessments of costs and benefits of security safeguards and consequences would need to be done twice: once in terms of readiness and once in terms of dollars. The safeguard efficacy assessments and initial probabilities are unaffected, however. Once the analysis cycle is finished, the decision maker would be presented with two separate analyses that together illustrate the amount of readiness purchased by a dollar of security.

This approach is often preferable to monetizing difficult-to-quantify values because calibration of such factors tends to vary considerably among stakeholders. The potential for disagreement and gridlock is greater if these assumptions are fixed at the beginning of the process when stakeholders cannot see the immediate impact of their calibration decisions. For this reason, the example analysis presented in this chapter purposely omitted them in an attempt to develop a model that could predict actual losses accurately. The example analysis also neglected to internalize risk attitude for the same reason. By grounding the analysis thus, a context is set for discussing intangible values and risk attitude. For example, if a decision maker were to choose a basket of safeguards that yielded a sub-optimal or even negative net benefit, then the difference between the optimal and chosen policies' net benefits could be considered the premium that the decision maker would be willing to pay for added security and the betterment of an intangible value. Upon further reflection, or perhaps after the intervention of stakeholders with different priorities, the decision maker might determine that the benefits of the additional security measures do not warrant their cost. In this way, the model can provide a solid basis for fruitful and informed discussions and decisions about intangible values and risk attitude.

This chapter has illustrated by example the process of developing an analytical model for managing computer security risks by marrying the proposed methodology to anecdotal data from the CSI/FBI survey. The subsequent model analysis demonstrates how uncertain data and expert judgments can be used to derive powerful results that not only direct future modeling and information gathering activities but also focus attention on the material managerial issues. The model's extensibility, adaptability, and applicability to multiple organizations with varying security management needs make it an important and valuable tool for analyzing and facilitating risk-management decision making.

Chapter 5 Conclusions and Implications

He said that there was one only good, namely, knowledge; and one only evil, namely, ignorance.

— Diogenes Laërtius, *The Lives and Opinions of Eminent Philosophers*⁸³

In this dissertation, I have traced the evolution of computer security risk modeling from its early days to its present incarnation, and predicted where I believe current trends will take it. Initial research in the late 1980s introduced the concept of Annual Loss Expectancy (ALE) as a risk metric and laid the foundation for subsequent modeling efforts. The failure of these initial ALE-based methodologies to resolve problems of implementation complexity, dearth of data, and issues of organizational acceptance led to the emergence of a second generation of modeling approaches, which were distinguished by their distinct detachment from data. Although meeting with some measure of success, this second generation is but a transitory step to a new generation, the seeds of which are being planted as insurance companies begin to assume a role in the management of computer risks.

The decision-analysis-based, ALE framework that I propose addresses several of the challenges facing quantitative modeling efforts by explicitly incorporating uncertainty, flexibly allowing for varying levels of modeling detail, placing the modeling focus squarely on the management decision, and recognizing the importance of computer security statistics. The forced explication of underlying assumptions about key quantities in a risk assessment provides a context for understanding the decision alternatives and the biases of the people involved. The adaptability and extensibility of the modeling approach make it generically applicable to virtually any computer security risk-management decision. The tools of decision analysis can be adroitly applied in a process of progressive refinement to balance model fidelity with tractability.

The analysis of publicly available data found the data to be woefully inadequate for supporting computer security risk-management decisions. The best available data is only anecdotal and not representative of any specific industry or group. The need for improved data sharing, collecting, and standardizing remains as pressing today, if not even more so, as it was in years past.

The suggested approach for valuing the consequences of computer security breaches calls for a counterfactual exercise that examines how an organization's future is materially affected by a security breach and compares that future to a second assessment of how the organization would have fared without the security breach. This concept represents a marked departure from the retrospective accounting techniques often employed when values are attached to information assets and to consequences of breaches in computer security.

Finally, the case study example demonstrates how uncertain data and expert judgments can be combined in the proposed decision-analysis framework to inform decisions about computer security risk management. The model analysis shows the relative importance of different input variables and assumptions, the value of additional information and where future model efforts should be focused, and the risk trade-offs between competing policies. Using publicly available, anecdotal data, the model showed quite convincingly that the

⁸³Diogenes Laërtius, *The Lives and Opinions of Eminent Philosophers*, translated by Charles Duke Yonge (London: George Bell & Sons, 1895), Chapter XIV.

current level of reported computer-security-related risks warrants only the most inexpensive of additional safeguards. Unless the costs and consequences of computer security breaches are radically erroneous, the optimal solution for managing computer security risks calls for very minimal security measures. Thus, the reluctance of both private and government organizations to pursue computer security aggressively may be well justified. Of course, this conclusion is very weak because it rests on an application of anecdotal data that many security experts agree underestimate the true extent and consequences of computer crime.

The implications of this research for government policy are consistent with the findings of several organizations, both government-sponsored and private, with respect to computer security incident information sharing.⁸⁴ The need for higher quality, more comprehensive information sharing to elucidate both the frequency with which computer security incidents are occurring and the severity of their consequences is essential to any attempt at assessing the robustness and security of the national infrastructure. To the extent that government can facilitate this process by easing antitrust restrictions on company interactions and clarifying the limits of legal liability incurred by organizations during their stewardship of others' private information, legal obstacles to information sharing may be lessened or removed. The question of whether the government must assume an additional role of actively gathering, analyzing, and disseminating such information because private organizations cannot or will not rise to the challenge remains to be seen. Regardless, the need to ascertain a more accurate quantitative picture of the country's infrastructure security posture, potential vulnerability to computer-related attacks, and overall risk is of pivotal importance not only to security resource allocation decisions but to the stability of an infrastructure that plays a large and growing role in the economy and society.

Future research in this area would do well to develop tools and metrics to aid in the gathering of information security statistics, specifically statistics that support risk-management activities. Nascent efforts are already under way in technical computer security circles to standardize the jargon and lend more rigorous definition to its terms.⁸⁵ Standards for measuring losses, technology for detecting incidents, and automated tools for dealing with both are all areas pregnant with difficult, researchable problems. Linkages will also be needed to extend the modeling framework to new network modeling and simulation tools that are currently under development. These tools will simulate computer and network environments so as to observe incident patterns and attacker behavior. These simulations could be used to derive significantly more accurate measures of safeguard efficacy and to ascertain the network-relevant consequences of security incidents. Another extension of this research might examine the gaming aspect of computer security. As attacker behavior becomes better understood, economic game theory may be able to predict how various attackers and defenders will behave under different security policies. Thus, better assessments of both security policy effectiveness and likelihoods of successful attacks may be achievable.

Computer security risk management today finds itself unable to answer satisfactorily the question, "How much is enough?" With the second-generation methodologies firmly entrenched, only limited movements are being made to address this most basic of questions. The return to a quantitative approach, like the one presented in this

⁸⁴ For example, see the Center for International Security and Cooperation Workshop series on Protecting and Assuring Critical Infrastructures, March and July 1997; or the President's Commission on Critical Infrastructure Protection, *Critical Foundations: Protecting America's Infrastructures* (Washington, DC: U.S. Government Printing Office, October 1997).

⁸⁵ See John D. Howard and Thomas A. Longstaff, *A Common Language for Computer Security Incidents*, SAND98-8667 (Albuquerque: Sandia National Laboratories, October 1998).

dissertation, is inevitable. Forces such as insurance, legal liability, and market competition will only expedite this process. As society's dependence upon digital computing and telecommunications increases, the need for quantitative computer security risk management will become more acute. Eventually, computer security risk management will be compelled to abandon its folk-art ways, make Bernstein's rite of passage to the modern era, and assume its rightful place alongside other once-inscrutable risks that are now actively and effectively managed.

Appendix

The following software code can be read from a plain text, or ASCII, file by evaluation versions of Analytica 1.1.1 or later on either the Macintosh or Windows platform. To download a free evaluation copy of Analytica from Lumina Decision Systems, Inc., go to its web site at <<http://www.lumina.com>>. The evaluation version will enable the user to tour the model, manipulate values, and see results. Model construction and extension will require the purchase of a full professional version of Analytica from Lumina Decision Systems, Inc.

```
{ From user Kevin J. Soo Hoo, Model Example_model at Thu, Mar 9, 2000 1:14 AM~~
}
Softwareversion 1.1.1
```

```
{ System Variables with non-default values: }
Windows := 2
Heapsize := 6.443M
Typechecking := 1
Checking := 1
Graphwindows := 5
Showundef := 0
Saveoptions := 2
Savevalues := 0
Webhelper := -1
```

```
{ Non-default Time SysVar value: }
Time := [0,1,2]
Title Time: Time
```

```
Model Example_model
Title: Computer Security Risk Management Model
Description: This model describes the basic decision variables that fo~~
rm the basis any computer security resource allocation decision. By e~~
ntering appropriate data, the user may see a cost-benefit trade-off bet~~
ween between different security policies.
Author: Kevin J. Soo Hoo
Date: Thu, Nov 4, 1999 1:03 AM
Saveauthor: Kevin J. Soo Hoo
Savedate: Thu, Mar 9, 2000 1:14 AM
Defaultsize: 48,24
Diagstate: 1,46,42,667,449,17
Fileinfo: 0,Model Example_model,1,1,0,Eledhel:Stanford:Dissertation:Ch~~
apterhouse:Example Model
Pagesetup: (00030000004800480000000002D80228FFE1FFE202F902460347052803~~
```

FC00020000004800480000000002D8022800010000006400000001000303030000000
 1~~
 270F000100010000000000000000000000006008001901900000000000000000000000
 ~~
 000000000000000000000000000000000000000000000000000000000000000000000)

Text Text1
 Description: Computer Security Risk Management Decision
 Nodelocation: 344,32
 Nodesize: 236,12
 Nodeinfo: 1,0,0,1,0,0,1,,0,
 Nodefont: Times, 24

Module Model_diagram
 Title: Model Diagram~
 (Double-Click Here to See)
 Description: This module holds the model details in a decision diagram~~
 format.
 Author: Kevin Soo Hoo
 Date: Fri, Nov 12, 1999 11:22 AM
 Defaultsize: 48,24
 Nodelocation: 472,240
 Nodesize: 52,24
 Diagstate: 1,356,45,390,272,17

Decision Safeguard_selector
 Title: Safeguard Selector
 Description: This variable is used to select the safeguards to be incl~~
 uded in each policy option. A 1 represents inclusion and a 0 represen~~
 ts exclusion.
 Definition: Table(Policy_index,Safeguards)(
 0,0,0,0,0,0,0,0,0,0,0,0,
 0,1,0,1,0,0,0,1,0,0,1,0,
 1,1,0,1,1,0,1,1,1,0,1,0,
 1,1,1,1,1,1,1,1,1,1,1,
)
 Nodelocation: 64,104
 Nodesize: 44,20
 Nodeinfo: 1,1,1,1,1,0,,1,
 Defnstate: 1,241,466,750,302,0,MIDM
 Aliases: Alias Safeguard_selector1, Formnode Policies__groups_of_
 Reformdef: [Policy_index, Safeguards]

Chance Costs_of_safeguards
 Title: Costs of Safeguards
 Units: \$
 Description: Annual Costs associated with safeguard implementation. T~~
 hese costs were culled from industry reports, price quotations, and su~~
 rvey data.
 Definition: Table(Safeguards)(
 Uniform(200K,600K),Uniform(500K,600K),Uniform(600K,1M),560K,Uniform(1M~~

,1.1M),Uniform(8M,12M),Uniform(288K,420K),10K,Uniform(900K,1.3M),Unifor
rm(65K,70K),Uniform(75K,150K),Uniform(200K,300K))
Nodelocation: 64,48
Nodesize: 44,20
Nodeinfo: 1,1,1,1,1,1,0,,1,
Windstate: 1,184,327
Defnstate: 1,606,84,355,333,0,MIDM
Aliases: Formnode Costs_of_safeguards1

Variable Policy_cost
Title: Added Costs
Units: \$
Description: Costs of each safeguard are added together to compute the
annual cost of each policy option.
Definition: Sum((Costs_of_safeguards*Safeguard_selector),Safeguards)
Nodelocation: 192,48
Nodesize: 48,24

Variable New_revenues_enabled
Title: Add'l Profits enabled by Safeguards
Units: \$
Description: The adoption of some safeguards will enable new sources o
f revenue. To the extent that the profits from these revenues can be
captured in the analysis, they should be included. For the example an
alysis, this factor has been ignored.
Definition: Table(Safeguards)(
0,0,0,0,0,0,0,0,0,0,0,0)
Nodelocation: 64,160
Nodesize: 44,24
Nodeinfo: 1,1,1,1,1,1,0,,1,
Aliases: Formnode New_profits_enabled_

Variable Policy_new_revenues
Title: Added Profits
Units: \$
Description: New profits associated with each policy option are taken
and grouped together into the profits associated with security policie
s.
Definition: Sum((New_revenues_enabled*Safeguard_selector),Safeguards)
Nodelocation: 192,160
Nodesize: 48,24

Objective Cost_benefit_calcula
Title: Net Benefit
Units: \$/year
Description: Net Benefit of adopting each policy option. This value i
s calculated relative to the status quo.
Definition: (Annual_loss_expectan[policy_index='Status Quo']-Annual_lo
ss_expectan)+Policy_new_revenues-Policy_cost
Nodelocation: 312,104
Nodesize: 44,20

Valuestate: 1,166,129,654,413,0,MEAN
Aliases: Formnode Policy_cost_benefit_
Graphsetup: Graphtool:0~
Distresol:70~
Diststeps:1~
Cdfresol:2~
Cdfsteps:1~
Symbolsize:6~
Linestyle:1~
Frame:1~
Grid:1~
Ticks:1~
Mesh:1~
Scales:1~
Rotation:45~
Tilt:0~
Depth:70~
Frameauto:1~
Showkey:1~
Xminimum:-20M~
Xmaximum:5M~
Yminimum:0~
Ymaximum:1~
Zminimum:1~
Zmaximum:4~
Xintervals:0~
Yintervals:0~
Includexzero:0~
Includeyzero:0~
Includezzero:0~
Statsselect:[1,1,1,1,1,0,0,0]~
Probindex:[0.05,0.25,0.5,0.75,0.95]~

Index Policy_index

Title: Policy Index

Description: The four security policy options under consideration.

Definition: ['Status Quo','Minimal Improvement','Major Improvement','Maximum Improvement']

Nodelocation: 64,216

Nodesize: 44,16

Index Safeguards

Title: Safeguards

Description: Listing of all safeguards that could be included in the policy options. These safeguards are all new ones that might be adopted by the organization to improve its security.

Definition: ['Security Awareness','HW/SW Network Upgrade','Response Team','Nightly Back-ups','Encryption','Central Access Control','Firewalls','Screen Locking Software','Security Management Team','Comm Content Screening','Anti-Virus Software','Intrusion Detection System']

Nodelocation: 192,216
Nodesize: 44,16
Windstate: 1,408,418
Aliases: Formnode Safeguards1

Index Bad_events
Title: Bad Events
Description: Listing of undesired events that could result from computer security breaches.
Definition: ['Information Theft','Information Modification','Information Destruction','System Outage','Employee Theft','System Degradation']

Nodelocation: 312,216
Nodesize: 44,16
Aliases: Formnode Bad_events1

Module Annual_loss_module
Title: Annual Loss Module
Description: This module computes the annual loss expectancy for each policy option.
Author: Kevin J. Soo Hoo
Date: Mon, Jan 3, 2000 1:33 PM
Nodelocation: 192,104
Nodesize: 44,20
Diagstate: 1,359,78,391,279,17

Chance Initial_frequency_of
Title: Initial Frequency of Bad Events
Units: /year
Description: Initial assumption or data on the average number of bad events that occur in a year. ~
~
These values were taken from the "1999 CSI/FBI Computer Crime and Security Survey," Computer Security Issues and Trends, Vol. 5, No. 1, Computer Security Institute, Winter 1999. The distributions were formed using 1997, 1998, and 1999 incident data for "Theft of Proprietary Information, System Penetration by Insiders, Sabotage, Denial of Service, and Financial Fraud, Virus."
Definition: Table(Bad_events)(
Triangular((82/458),(101/492),(104/405)),Triangular((198/492),(203/458),(223/405)),Triangular((53/405),(69/492),(66/458)),Uniform((114/458),(129/405)),Triangular((59/492),(58/405),(68/458)),Triangular((407/492),(380/458),(365/405)))
Nodelocation: 72,88
Nodesize: 44,20
Nodeinfo: 1,1,1,1,1,1,0,,1,
Windstate: 1,146,261
Defnstate: 1,297,306,502,223,0,MIDM
Valuestate: 1,256,176,577,415,0,STAT
Aliases: Formnode Initial_frequency_o1
Graphsetup: Graphtool:0~

Distresol:10~
 Diststeps:1~
 Cdfresol:5~
 Cdfsteps:1~
 Symbolsize:6~
 Linestyle:1~
 Frame:1~
 Grid:1~
 Ticks:1~
 Mesh:1~
 Scales:1~
 Rotation:45~
 Tilt:0~
 Depth:70~
 Frameauto:1~
 Showkey:1~
 Xminimum:0~
 Xmaximum:1~
 Yminimum:0~
 Ymaximum:200~
 Zminimum:1~
 Zmaximum:6~
 Xintervals:0~
 Yintervals:0~
 Includexzero:0~
 Includeyzero:0~
 Includezzero:0~
 Statsselect:[1,1,1,1,1,0,0,0]~
 Probindex:[0.05,0.25,0.5,0.75,0.95]~

Variable Efficacy_of_safeguar
 Title: Reduction in Frequency
 Units: %
 Description: Percentage reduction in the frequencies of bad events that the adoption of a safeguard will precipitate. These numbers are based on expert elicitations.
 Definition: Table(Safeguards,Bad_events)(
 0.35,0.3,0.3,0.05,0.6,0.5,
 0.45,0.45,0.45,0.45,0,0.45,
 0.4,0.4,0.4,0.4,0,0.2,
 0,0,0,0,0,0,
 0,0,0,0,0,0,
 0.3,0.15,0.15,0,0.5,0,
 0.75,0.75,0.75,0.75,0.2,0.1,
 0.15,0.2,0.2,0,0.4,0,
 0.5,0.5,0.5,0.5,0.5,0.5,
 0.75,0,0,0,0.3,0,
 0,0.35,0.4,0,0,0.4,
 0.51,0.51,0.51,0.51,0.25,0.51
)

Nodelocation: 72,40
Nodesize: 44,20
Nodeinfo: 1,1,1,1,1,1,0,,1,
Defnstate: 1,52,218,956,301,0,MIDM
Aliases: Formnode Reduction_in_frequen
Reformdef: [Bad_events, Safeguards]

Variable Consequences_of_bad_
Title: Consequences of Bad Events
Units: \$

Description: This variable computes the consequences of bad events based on the safeguards selected in each policy, the reduction in consequences that those safeguards are expected to have, and the initial estimate of the consequences of bad events.

Definition: (Initial_consequences*Product((1-Safeguard_selector*Reduction_in_consequ),Safeguards))

Nodelocation: 200,208
Nodesize: 44,20
Valuestate: 1,40,50,692,244,0,MEAN
Reformval: [Policy_index, Bad_events]

Variable Annual_loss_expectan

Title: Annual Loss Expectancy

Units: \$/year

Description: Annual expected loss.

Definition: Sum((Consequences_of_bad_*frequency_of_bad_events),Bad_events)

Nodelocation: 312,136
Nodesize: 44,20
Valuestate: 1,125,49,496,309,0,MEAN
Aliases: Formnode Annual_loss_expecta1
Graphsetup: Graphtool:0~
Distresol:200~
Diststeps:1~
Cdfresol:5~
Cdfsteps:1~
Symbolsize:6~
Linestyle:1~
Frame:1~
Grid:1~
Ticks:1~
Mesh:1~
Scales:1~
Rotation:45~
Tilt:0~
Depth:70~
Frameauto:1~
Showkey:1~
Xminimum:-1M~
Xmaximum:3M~
Yminimum:0~

Ymaximum:5u~
 Zminimum:1~
 Zmaximum:4~
 Xintervals:0~
 Yintervals:0~
 Includexzero:0~
 Includeyzero:0~
 Includezzero:0~
 Statsselect:[1,1,1,1,1,0,0,0]~
 Probinde: [0.05,0.25,0.5,0.75,0.95]~

Variable Reduction_in_consequ
 Title: Reduction in Consequence
 Units: %

Description: Degree to which safeguards, if adopted, can reduce the c~
 onsequences of a bad event. These values represent expert judgments.

Definition: Table(Bad_events,Safeguards)(
 0,0,0,0,0.95,0,0,0,0,0,0,
 0,0,0,0.6,0.95,0,0,0,0,0,0,
 0,0,0,0.95,0,0,0,0,0,0,0,
 0,0,0.7,0,0,0,0,0,0,0,0,
 0,0,0,0,0,0,0,0,0,0,0,
 0,0,0.65,0,0,0,0,0,0,0,0

)
 Nodelocation: 72,184
 Nodesize: 44,20
 Nodeinfo: 1,1,1,1,1,1,0,,1,
 Defnstate: 1,120,101,945,368,0,MIDM
 Aliases: Formnode Reduction_in_conseq1
 Reformdef: [Bad_events, Safeguards]

Chance Initial_consequences
 Title: Initial Consequences of Bad Events
 Units: \$

Description: Probabilistic distributions of the financial damage that ~
 a bad event will cause.~

~
 Values taken from "1999 CSI/FBI Computer Crime and Security Survey," C~
 omputer Security Issues and Trends, Vol. 5, No. 1, Computer Security ~
 Institute, Winter 1999. The distributions were formed using the 1999 ~
 highest, average, and lowest reported losses for "Theft of Proprietar~
 y Information, Virus, Sabotage, Denial of Service, and Financial Fraud~
 , System Penetration by Outsider."

Definition: Table(Bad_events)(
 Triangular(1000,1.847652M,25M),Triangular(1000,103.142K,500K),Triangul~
 ar(1000,163.74K,1M),Triangular(1000,116.25K,1M),Triangular(10K,1.47059~
 2M,20M),Triangular(1000,45.465K,1M))
 Nodelocation: 72,232
 Nodesize: 44,20
 Nodeinfo: 1,1,1,1,1,1,0,,1,

Defnstate: 1,141,265,416,303,0,MIDM
Aliases: Formnode Initial_consequence1

Variable Frequency_of_bad_eve
Title: Frequency of Bad Events
Units: /year

Description: This variable computes the frequency of bad events based upon the safeguards selected for each policy, the initial frequency of bad events estimate, and the effect that selected safeguards will have on that initial frequency.

Definition: $\text{Initial_frequency_of} * \text{product}(1 - \text{Efficacy_of_safeguard_selector}, \text{safeguards})$

Nodelocation: 200,72

Nodesize: 44,20

Valuestate: 1,120,130,731,421,0,MEAN

Reformval: [Policy_index, Bad_events]

Alias Safeguard_selector1

Title: Safeguard Selector

Definition: 1

Nodelocation: 72,136

Nodesize: 44,20

Nodeinfo: 1,1,1,1,1,1,0,,1,

Original: Safeguard_selector

Close Annual_loss_module

Close Model_diagram

Formnode Initial_frequency_o1

Title: Initial Frequency of Bad Events

Definition: 0

Nodelocation: 160,136

Nodesize: 152,20

Original: Initial_frequency_of

Formnode Reduction_in_frequen

Title: Reduction in Frequency

Definition: 0

Nodelocation: 160,112

Nodesize: 152,20

Original: Efficacy_of_safeguard

Formnode Costs_of_safeguards1

Title: Costs of Safeguards

Definition: 0

Nodelocation: 160,272

Nodesize: 152,20

Original: Costs_of_safeguards

Formnode New_profits_enabled_

Title: New Profits enabled by Safeguards
Definition: 0
Nodelocation: 160,296
Nodesize: 152,24
Original: New_revenues_enabled

Formnode Safeguards1
Title: Safeguards
Definition: 0
Nodelocation: 160,376
Nodesize: 152,20
Original: Safeguards

Formnode Bad_events1
Title: Bad Events
Definition: 0
Nodelocation: 160,400
Nodesize: 152,20
Original: Bad_events

Formnode Reduction_in_conseq1
Title: Reduction in Consequence
Definition: 0
Nodelocation: 160,160
Nodesize: 152,20
Original: Reduction_in_consequ

Formnode Initial_consequence1
Title: Initial Consequences of Bad Events
Definition: 0
Nodelocation: 160,248
Nodesize: 152,20
Original: Initial_consequences

Formnode Policies__groups_of_
Title: Policies (Groups of Safeguards)
Definition: 0
Nodelocation: 480,112
Nodesize: 152,24
Original: Safeguard_selector

Formnode Policy_cost_benefit_
Title: Policy Cost/Benefit Calculation
Definition: 1
Nodelocation: 488,400
Nodesize: 156,24
Original: Cost_benefit_calcula

Formnode Annual_loss_expecta1
Title: Annual Loss Expectancy
Definition: 1

Nodelocation: 488,376
Nodesize: 156,20
Original: Annual_loss_expectan

Text Text2
Description: Statistics
Nodelocation: 152,80
Nodesize: 48,12
Nodeinfo: 1,0,0,1,0,0,1,,0,
Nodefont: Geneva, 18

Text Text3
Title: Text2
Description: Financial Estimates
Nodelocation: 160,216
Nodesize: 92,12
Nodeinfo: 1,0,0,1,0,0,1,,0,
Nodefont: Geneva, 18

Text Text4
Title: Text2
Description: Lists
Nodelocation: 144,344
Nodesize: 28,12
Nodeinfo: 1,0,0,1,0,0,1,,0,
Nodefont: Geneva, 18

Text Text5
Title: Text2
Description: Results
Nodelocation: 488,344
Nodesize: 40,12
Nodeinfo: 1,0,0,1,0,0,1,,0,
Nodefont: Geneva, 18

Text Text6
Title: Text2
Description: Policy Selector
Nodelocation: 472,80
Nodesize: 76,12
Nodeinfo: 1,0,0,1,0,0,1,,0,
Nodefont: Geneva, 18

Close Example_model

Bibliography

- Belleville, Patrice, "Report on the 2nd Risk Model Builders Workshop," *Proceedings of 1989 Computer Security Risk Management Model Builders Workshop*, National Institute of Standards and Technology, June 1989.
- Bellovin, Steven M., "Security Problems in the TCP/IP Protocol Suite," *Computer Communication Review*, Vol. 19, No. 2, April 1989, pp. 32–48.
- , "There Be Dragons," *Proceedings of the Third Usenix UNIX Security Symposium*, Baltimore, September 1992.
- Bernstein, Peter L., *Against the Gods: The Remarkable Story of Risk* (New York: John Wiley & Sons, Inc., 1996).
- Brown, Rex V., "Personalized Decision Analysis As an Expert Elicitation Tool: An Instructive Experience in Information Security Policy," Report submitted to the Office of Technology Assessment, February 25, 1985.
- Browne, Peter S. and Laverly, James E., "Using Decision Analysis to Estimate Computer Security Risk," *Proceedings of 1988 Computer Security Risk Management Model Builders Workshop*, National Bureau of Standards, May 1988, pp. 117–134.
- BSI/DISC Committee BDD/2, *British Standard 7799* (London: BSI, 1999).
- "The Buck Stops Here," *Information Security*, July 1999, <<http://www.infosecuritymag.com/july99/buck.htm>> last accessed March 10, 2000.
- "Budgets & Product Purchasing Trends," *Information Security*, July 1999, <<http://www.infosecuritymag.com/july99/chart2.htm>> last accessed March 10, 2000.
- Carroll, John M., *Managing Risk: A Computer-Aided Strategy* (Boston: Butterworth Publishers, 1984).
- Center for International Security and Cooperation Workshop series on Protecting and Assuring Critical Infrastructures, March and July 1997.
- Cheswick, William, "How Computer Security Works," *Scientific American*, Vol. 279, No. 4, October 1998, pp. 106–109.
- , *An Evening with Berferd in Which a Cracker Is Lured, Endured, and Studied*, white paper, AT&T Bell Laboratories.
- CIC Security Working Group, *Incident Cost Analysis and Modeling Project* (Ann Arbor, MI: University of Michigan, 1998).
- Cohen, Fred, "Managing Network Security: Balancing Risk," December 1998, <<http://all.net/journal/netsec/9812.html>> last accessed March 10, 2000.
- Cohen, Fred, Phillips, Cynthia, et al., "A Preliminary Classification Scheme for Information System Threats, Attacks, and Defenses; A Cause and Effect Model; and Some Analysis Based on that Model" (Livermore, CA: Sandia National Laboratories, September 1998).

- Cooper, Robert G., Edgett, Scott J., and Kleinschmidt, Elko J., *Portfolio Management for New Products* (Reading, MA: Addison-Wesley, 1998).
- Corresponding Committee for Information Valuation, *Guideline for Information Valuation*, Will Ozier, editor (Glenview, IL: Information Systems Security Association, Inc., 1993).
- Dalton, Gregory, "Acceptable Risks," *Information Week*, August 31, 1998, <<http://www.informationweek.com/698/98iursk.htm>> last accessed March 10, 2000.
- Drake, David L., and Morse, Katherine L., "The Security-Specific Eight-Stage Risk Assessment Methodology," *Proceedings of the 17th National Computer Security Conference*, 1994.
- Economist Intelligence Unit, *Managing Business Risks in the Information Age* (New York: The Economist Intelligence Unit, Ltd., 1998).
- "Enough Is (Never) Enough," *Information Security*, July 1999, <<http://www.infosecuritymag.com/july99/enough.htm>> last accessed March 10, 2000.
- Epstein, Richard A., *Cases and Materials on Torts*, 6th edition (Boston: Little, Brown & Co., 1995).
- Feist, Raymond E., *Prince of the Blood* (New York: Doubleday, 1989).
- Friedman, Lawrence M., *A History of American Law*, 2nd edition (New York: Simon & Schuster, 1985).
- General Accounting Office, *Information Security Risk Assessment: Practices of Leading Organizations, Exposure Draft* <<http://www.gao.gov/monthly.list/aug99/aug991.htm>> last accessed 1 October 1999.
- Gibbons, John H., *Cybernation: The American Infrastructure in the Information Age* <<http://www.whitehouse.gov/WH/EOP/OSTP/html/cyber2.html>> last accessed March 10, 2000.
- Gilbert, Irene E., "Guide for Selecting Automated Risk Analysis Tools," National Institute of Standards and Technology, Special Publication #SP 500-174, October 1989.
- "Got Security?" *Information Security Magazine*, July 1999, <<http://www.infosecuritymag.com/july99/cover.htm>> last accessed March 10, 2000.
- GSSP Draft Sub-committee, *Exposure Draft of the Generally Accepted System Security Principles (GSSP)*, Datapro Information Services, August 18, 1994.
- Guarro, Sergio B. "Principles and Procedures of the LRAM Approach to Information Systems Risk Analysis and Management," *Computers and Security*, Vol. 6, Elsevier Science Publishers B.V., 1987.
- Guarro, Sergio B., "Analytical and Decision Models of the Livermore Risk Analysis Methodology (LRAM)," *Proceedings of 1988 Computer Security Risk Management Model Builders Workshop*, National Bureau of Standards, May 1988, pp. 49–72.

- Hansell, Saul, "A \$10 Million Lesson in the Risks of Electronic Banking (Hacker Taps into Citibank Accounts)," *New York Times*, vol. 144, August 19, 1995, pp. 15(N), 31(L) col. 2.
- Harvard Business Review on Managing Uncertainty* (Boston: Harvard Business School Press, 1999).
- Helsing, Cheryl, Swanson, Marianne, and Todd, Mary Anne, *Executive Guide to the Protection of Information Resources* (Gaithersburg, MD: National Institute of Standards and Technology, 1988).
- Hoffman, Lance J. and Hung, Brian T., "A Pictorial Representation and Validation of the Emerging Computer System Security Risk Management Framework," *Proceedings of the Computer Security Risk Management Model Builders Workshop*, Ottawa, Canada, June 20–22, 1989.
- Hoffman, Lance J., "A Prototype Implementation of a General Risk Model," *Proceedings of 1988 Computer Security Risk Management Model Builders Workshop*, National Bureau of Standards, May 1988, pp. 135–144.
- Hooper, T. J., 60 F.2d 737 (2d Cir. 1932).
- Howard, John D. and Longstaff, Thomas A., *A Common Language for Computer Security Incidents*, SAND98-8667 (Albuquerque: Sandia National Laboratories, October 1998).
- Howard, John D., *An Analysis of Security Incidents on the Internet 1989–1995*, Ph.D. thesis, Department of Engineering and Public Policy, Carnegie Mellon University, April 7, 1997.
- Howard, Ronald A. and Matheson, James E., "Influence Diagrams," *Readings in the Principles and Applications of Decision Analysis*, Vol. 2, Ronald A. Howard & James E. Matheson, editors (Menlo Park, CA: Navigant Consulting, Inc., 1983).
- Howard, Ronald A., "Decision Analysis: Applied Decision Theory," *Proceedings of the Fourth International Conference on Operational Research*, David B. Hertz and Jacques Melese, editors (New York: Wiley-Interscience, 1966).
- Hutt, Arthur E., Bosworth, Seymour and Hoyt, Douglas B., *Computer Security Handbook* (New York: John Wiley & Sons, Inc., 1995).
- Jacobson, Robert V., "IST/AMP and CRITI-CALC: Risk Management Tools," *Proceedings of 1988 Computer Security Risk Management Model Builders Workshop*, National Bureau of Standards, May 1988, pp. 73–88.
- Jansma, Roxana M., Sharon K. Fletcher, et al., *Risk-Based Assessment of the Surety of Information Systems* (Springfield, VA: National Technical Information Service, 1996).
- Jarworski, Lisa M., "Tandem Threat Scenarios: A Risk Assessment Approach," *Proceedings of the 16th National Computer Security Conference*, 1993.
- Jenkins, B. D., *Security Risk Analysis and Management*, white paper, Countermeasures Corporation, <<http://www.buddysystem.net/download/budsyswp.pdf>> last accessed March 10, 2000.
- Johnson, Samuel, *The History of Rasselas, Prince of Abyssinia: A Tale* (London: Printed for John Sharpe, 1817).

- Joint Security Commission, *Redefining Security; A Report to the Secretary of Defense and the Director of Central Intelligence* (Washington, DC: U.S. GPO) February 28, 1994.
- Katzke, Stuart W., "A Government Perspective on Risk Management of Automated Information Systems," *Proceedings of 1988 Computer Security Risk Management Model Builders Workshop*, National Bureau of Standards, May 1988, pp. 3–20.
- Kumamoto, Hiromitsu, and Henley, Ernest J., *Probabilistic Risk Assessment and Management for Engineers and Scientists*, 2nd edition (New York: Institute of Electrical and Electronics Engineers, Inc., 1996).
- Laërtius, Diogenes, *The Lives and Opinions of Eminent Philosophers*, translated by Charles Duke Yonge (London: George Bell & Sons, 1895), Chapter XIV.
- Larsen, Amy K., "Global Security Survey: Virus Attack," *Information Week*, July 12, 1999 <<http://www.informationweek.com/743/security.htm>> last accessed March 10, 2000.
- , "Worldwide Security Priorities," *Information Week*, July 12, 1999 <<http://www.informationweek.com/743/securit3.htm>> last accessed March 10, 2000.
- Lavine, Charles H., "The Aerospace Risk Evaluation System (ARiES): Implementation of a Quantitative Risk Analysis Methodology for Critical Systems," *NIST/NCSC: 17th National Computer Security Conference*, Baltimore, October 11–14, 1994, pp. 431–440.
- Levy, Haim, *Stochastic Dominance: Investment Decision Making under Uncertainty* (Boston: Kluwer Academic Publishers, 1998).
- Levy, Steven, *Hackers: Heroes of the Computer Revolution* (New York: Dell Publishing, 1984).
- Matheson, David, and Matheson, Jim, *The Smart Organization: Creating Value through Strategic R&D* (Boston: Harvard Business School Press, 1998).
- Mayerfeld, Harold N., "Definition and Identification of Assets As the Basis for Risk Management," *Proceedings of 1988 Computer Security Risk Management Model Builders Workshop*, National Bureau of Standards, May 1988, pp. 21–34.
- Mayerfeld, Harold N., Troy, E. F., et al., "M²R_x: Model-Based Risk Assessment Expert," *AIAA/ASIS/IEEE 3rd Aerospace Computer Security Conference: Applying Technology to Systems*, Orlando, FL, December 7–11, 1987, pp. 87–92.
- Meinel, Carolyn P., "How Hackers Break In . . . and How They Are Caught," *Scientific American*, Vol. 279, No. 4, October 1998, pp. 98–105.
- Miora, Michael, "Quantifying the Business Impact Analysis," *Carolina Computer News*, September 1996.
- Mitre Corporation, *Common Vulnerabilities and Exposures*, <<http://www.cve.mitre.org>> last accessed March 10, 2000.
- Morgan, M. Granger, and Henrion, Max, *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*, 2nd edition (New York: Cambridge University Press, forthcoming).

- Morgan, M. Granger, and Henrion, Max, *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*, 1st edition (New York: Cambridge University Press, 1990).
- Moseleh, Ali, "A Matrix/Bayesian Approach to Risk Management of Information Systems," *Proceedings of 1988 Computer Security Risk Management Model Builders Workshop*, National Bureau of Standards, May 1988, pp. 103–116.
- National Bureau of Standards, *Guideline for Automatic Data Processing Risk Analysis*, FIPS PUB 65 (Washington, DC: U.S. General Printing Office, 1979).
- National Institute of Standards and Technology, *Description of Automated Risk Management Packages That NIST/NCSC Risk Management Research Laboratory Has Examined*, updated 1991 <<http://csrc.nist.gov/training/risktool.txt>> last accessed on 27 October 1999.
- National Institute of Standards and Technology, *Guideline for the Analysis of Local Area Network Security*, Federal Information Processing Standards Publication 191, FIPS PUB 191 (Gaithersburg, MD: National Institute of Standards and Technology, 1994).
- National Research Council, Committee on Risk Characterization, *Understanding Risk: Information Decisions in a Democratic Society*, Stern, Paul C., and Fineberg, Harvey V. (eds.) (Washington, DC: National Academy Press, 1996).
- National Research Council, *Trust in Cyberspace*, Fred B. Schneider, editor (Washington, DC: National Academy Press, 1999).
- NetSolve, *ProWatch Secure Network Security Survey* (Austin, TX: NetSolve Corporation, 1997).
- Neumann, Peter, speaking at the Conference on Cyber Crime and Terrorism, Hoover Institution, Stanford University, December 6, 1999.
- Ozier, Will, "Issues in Quantitative versus Qualitative Risk Analysis," Datapro Information Services, May 4, 1999.
- Palmer, I. C., and Potter, G. A., *Computer Security Risk Management* (London: Jessica Kingsley Publishers, 1989)
- Parker, Donn B., *Fighting Computer Crime: A New Framework for Protecting Information* (New York: John Wiley & Sons, Inc., 1998).
- Perspective, "Crime's Cost," *Investor's Business Daily*, May 9, 1996.
- Posner, Richard A., *Economic Analysis of Law*, 4th edition (Boston: Little, Brown & Co., 1992).
- Power, Richard, "1996 CSI/FBI Computer Crime and Security Survey," *Computer Security Issues & Trends*, Vol. 2, No. 2, Spring 1996.
- _____, "1998 CSI/FBI Computer Crime and Security Survey," *Computer Security Issues & Trends*, Vol. 4, No. 1, Winter 1998.
- _____, "1999 CSI/FBI Computer Crime and Security Survey," *Computer Security Issues & Trends*, Vol. 5, No. 1, Winter 1999.
- _____, Richard, "CSI Special Report: How to Quantify Financial Losses from Infosec Breaches?" *Computer Security Alert*, October 1999.

- President's Commission on Critical Infrastructure Protection, *Critical Foundations: Protecting America's Infrastructures* (Washington, DC: U.S. Government Printing Office, October 1997).
- Proceedings of the Computer Security Risk Management Model Builders Workshop* (Washington, DC: National Institute of Standards and Technology, 1988).
- Raiffa, Howard, and Schlaifer, Robert, *Applied Statistical Decision Theory* (Boston: Harvard University, 1961).
- Rivest, Ronald L., "The Case against Regulating Encryption Technology," *Scientific American*, Vol. 279, No. 4, October 1998, pp. 106–107.
- Russell, Deborah, and Gangemi, G. T. Sr., *Computer Security Basics* (New York: Thunder Mountain Press, 1994).
- Safeware, Inc., "Safeware's 1998 Loss Study," May 1, 1999, <<http://www.safeware.com/safeware/pressreleases.htm>> last accessed March 10, 2000.
- Schmidt, Edwin A., "Conceptual Model of the Risk Management Process," *Proceedings of 1988 Computer Security Risk Management Model Builders Workshop*, National Bureau of Standards, May 1988, pp. 89–102.
- "Security of the Internet," *The Froehlich/Kent Encyclopedia of Telecommunications*, Vol. 15 (New York: Marcel Dekker, 1997), pp. 231–255.
- Smith, Suzanne T., "LAVA: An Expert System Framework for Risk Analysis," *Proceedings of 1988 Computer Security Risk Management Model Builders Workshop*, National Bureau of Standards, May 1988, pp. 179–202.
- Smith, Suzanne T., and Jalbert, M. L., "LAVA/CIS Version 2.0: A Software System for Vulnerability and Risk Assessment," *NIST/NCSC 13th National Computer Security Conference*, Washington, DC, October 1–4, 1990, pp. 460–469.
- Snow, David W., "A General Model for the Risk Management of ADP Systems," *Proceedings of 1988 Computer Security Risk Management Model Builders Workshop*, National Bureau of Standards, May 1988, pp. 145–162.
- Staten, Clark L., *Results of ERRI/EmergencyNet News Local/Count/State Computer "Hacking" Survey—May/June, 1999*, July 19, 1999, <<http://www.emergency.com/1999/hackrslt.htm>> last accessed March 10, 2000.
- Strassman, Paul A., "Risk Management: The Need for a Shared View," presentation to The Conference Board, Information Technology Fluency Conference, New York, NY, June 15, 1999.
- "Survey Overview & Executive Summary," *Information Security*, July 1999, <<http://www.infosecuritymag.com/july99/chart1.htm>> last accessed March 10, 2000.
- Swartwood, Dan T., and Hefferman, Richard J., "ASIS Trends in Intellectual Property Loss Survey Report," (Alexandria, VA: American Society for Industrial Security, International, 1998).
- Titus vs. Bradford, B. & K. R. Co., 20 A. 517 (Pa. 1890).

- Tompkins, Frederick G., "A Systems Approach to Information Security Risk Management" <<http://www.ncsa.com/knowledge/research/97072402.htm>> last accessed on August 11, 1998.
- Townsend, Timothy J., *Security Adequacy Review Process and Technology*, Technical White Paper (Palo Alto, CA: Sun Microsystems, 1998).
- Trigeorgis, Lenos, *Real Options: Managerial Flexibility and Strategy in Resource Allocation*, 4th printing, 1999 (Cambridge: MIT Press, 1996).
- Troy, Eugene F., "Introduction and Statement of Purpose," *Proceedings of 1988 Computer Security Risk Management Model Builders Workshop*, National Bureau of Standards, May 1988, pp. 1-2.
- "Under Attack & Underprepared," *Information Security*, July 1999, <<http://www.infosecuritymag.com/july99/under.htm>> last accessed March 10, 2000.
- United States vs. Carroll Towing Company, 159 F.2d 169, 173 (2d Cir. 1947).
- U.S. General Accounting Office, *Information Security: Computer Attacks at Department of Defense Pose Increasing Risks* (Washington, DC: U.S. Government Printing Office, May 1996).
- U.S. General Accounting Office, *Information Security Risk Assessment: Practices of Leading Organizations*, exposure draft, GAO-AIMD-99-139 (Washington, DC: U.S. Government Printing Office, August 1999).
- U.S. General Accounting Office, *Information Security Serious Weaknesses Put State Department and FAA Operations at Risk*, GAO/T-AIMD-98-170 (Washington, DC: U.S. Government Printing Office, May 1998).
- West-Brown, Moira, and Kossakowski, Klaus-Peter, *International Infrastructure for Global Security Incident Response* (Pittsburgh, PA: Carnegie Mellon University, 1999).
- Weston, Rusty, "Security Survey Methodology," *Information Week*, July 12, 1999 <<http://www.informationweek.com/743/securit2.htm>> last accessed March 10, 2000.
- Wood, Charles C., *Best Practices in Internet Commerce Security* (Sausalito, CA: Baseline Software, Inc., 1998).
- _____, *Information Security Policies Made Easy* (Sausalito, CA: Baseline Software, Inc., 1998).
- _____, "Information Security Staffing Levels: Calculating the Standard of Due Care," Baseline Software, Inc., March 11, 1998.
- _____, "Using Information Security to Achieve Competitive Advantage," in *Proceedings of the 18th Annual Computer Security Institute Conference*, Miami, November 11-15, 1991.
- Wyss, Gregory D., Craft, Richard L., Funkhouser, Donald R., *The Use of Object-Oriented Analysis Methods in Surety Analysis*, SAND99-1242 (Albuquerque: Sandia National Laboratories, May 1999).

- Wyss, Gregory D., Daniel, Sharon L., et al., *Information Systems Vulnerability: A Systems Analysis Perspective*, SAND96-1541C (Albuquerque: Sandia National Laboratories, 1996).
- Wyss, Gregory D., Fletcher, Sharon K., et al., *Toward a Risk-Based Approach to the Assessment of the Surety of Information Systems*, SAND95-0501C (Albuquerque: Sandia National Laboratories, 1995).
- Yoran, Amit, and Hoffman, Lance C., "Role-based Risk Analysis," *Proceedings of the 20th National Information Systems Security Conference*, October 1997.
- Zimmerman, Philip R., "Cryptography for the Internet," *Scientific American*, Vol. 279, No. 4, October 1998, pp. 110–115.

**Selected Reports, Working Papers, and Reprints
of the Center for International Security and Cooperation,
Stanford University**

To order, call (650) 725-6488 or fax (650) 723-0089. Selected publications and a complete publications list are also available on the center's website: www.stanford.edu/group/CISAC/.

- Herbert L. Abrams. *Can the Nation Afford a Senior Citizen As President? The Age Factor in the 1996 Election and Beyond*. 1997.
- David Alderson, David Elliott, Gregory Grove, Timothy Halliday, Stephen Lukasik, and Seymour Goodman. *Workshop on Protecting and Assuring Critical National Infrastructure: Next Steps*. 1998.
- Andrei Baev, Matthew J. Von Bencke, David Bernstein, Jeffrey Lehrer, and Elaine Naugle. *American Ventures in Russia. Report of a Workshop on March 20-21, 1995, at Stanford University*. 1995.
- Michael Barletta. *The Military Nuclear Program in Brazil*. 1997.
- David Bernstein. *Software Projects in Russia: A Workshop Report*. 1996.
- David Bernstein, editor. *Cooperative Business Ventures between U.S. Companies and Russian Defense Enterprises*. 1997.
- David Bernstein. *Commercialization of Russian Technology in Cooperation with American Companies*. 1999.
- George Bunn. *The Nonproliferation Regime under Siege*. 1999.
- George Bunn and David Holloway. *Arms Control without Treaties? Rethinking U.S.-Russian Strategic Negotiations in Light of the Duma-Senate Slowdown in Treaty Approval*. 1998.
- Irina Bystrova. *The Formation of the Soviet Military-Industrial Complex*. 1996.
- Jor-Shan Choi. *A Regional Compact Approach for the Peaceful Use of Nuclear Energy—Case Study: East Asia*. 1997.
- David Darchiashvili and Nerses Mkrttchian. *Caucasus Working Papers*. 1997.
- John S. Earle and Ivan Komarov. *Measuring Defense Conversion in Russian Industry*. 1996.
- Lynn Eden and Daniel Pollack. *Ethnopolitics and Conflict Resolution*. 1995.
- David Elliot, Lawrence Greenberg, and Kevin Soo Hoo. *Strategic Information Warfare—A New Arena for Arms Control?* 1997.
- Steve Fetter. *Climate Change and the Transformation of World Energy Supply*. 1999.
- Geoffrey E. Forden. *The Airborne Laser: Shooting Down What's Going Up*. 1997.
- James E. Goodby. *Can Strategic Partners Be Nuclear Rivals?* (First in a series of lectures on “The U.S.–Russian Strategic Partnership: Premature or Overdue?”) 1997.
- James E. Goodby. *Loose Nukes: Security Issues on the U.S.–Russian Agenda* (Second in a series of lectures on “The U.S.–Russian Strategic Partnership: Premature or Overdue?”) 1997.
- James E. Goodby. *NATO Enlargement and an Undivided Europe* (Third in a series of lectures on “The U.S.–Russian Strategic Partnership: Premature or Overdue?”) 1997.
- James E. Goodby and Harold Feiveson (with a foreword by George Shultz and William Perry). *Ending the Threat of Nuclear Attack*. 1997.
- Seymour Goodman. *The Information Technologies and Defense: A Demand-Pull Assessment*. 1996.
- Seymour Goodman, Peter Wolcott, and Patrick Homer. *High-Performance Computing, National Security Applications, and Export Control Policy at the Close of the 20th Century*. 1998.
- Lawrence T. Greenberg, Seymour E. Goodman, and Kevin J. Soo Hoo. *Old Law for a New World? The Applicability of International Law to Information Warfare*. 1997.
- Gregory D. Grove. *The U.S. Military and Civil Infrastructure Protection: Restrictions and Discretion under the Posse Comitatus Act*. 1999.
- Yunpeng Hao. *China's Telecommunications: Present and Future*. 1997.
- John R. Harvey, Cameron Binkley, Adam Block, and Rick Burke. *A Common-Sense Approach to High-Technology Export Controls*. 1995.

- Hua Di. *China's Security Dilemma to the Year 2010*. 1997.
- Alastair Iain Johnston, W. K. H. Panofsky, Marco Di Capua, and Lewis R. Franklin. *The Cox Committee Report: An Assessment*. 1999.
- Leonid Kistersky. *New Dimensions of the International Security System after the Cold War*. 1996.
- Taira Koybaeva, editor. *Strategic Stability and U.S.-Russian Relations. Report of the Twelfth Protocol Meeting between the Center for International Security and Arms Control and the Committee of Scientists for Global Security*. 1999.
- Amos Kovacs. *The Uses and Nonuses of Intelligence*. 1996.
- Allan S. Krass. *The Costs, Risks, and Benefits of Arms Control*. 1996.
- Gail Lapidus and Renée de Nevers, eds. *Nationalism, Ethnic Identity, and Conflict Management in Russia Today*. 1995.
- Liu Suping. *A Verification Regime for Warhead Control*. January 2000.
- Stephen J. Lukasik et al. *Review of the National Information Systems Protection Plan Version 1.0 March 5, 1999 Draft*. 1999.
- Kenneth B. Malpass et al. *Workshop on Protecting and Assuring Critical National Infrastructure*. 1997.
- Michael M. May. *Rivalries Between Nuclear Power Projectors: Why the Lines Will Be Drawn Again*. 1996.
- Michael M. May, editor. *The Cox Committee Report: An Assessment*. 1999.
- Michael M. May. *The U.S. Enlargement Strategy and Nuclear Weapons*. 2000.
- Robert L. Rinne. *An Alternative Framework for the Control of Nuclear Materials*. 1999.
- Vadim Rubin. *The Geopolitics of Energy Development in the Caspian Region: Regional or Conflict?* 1999.
- Xiangli Sun. *Implications of a Comprehensive Test Ban for China's Security Policy*. 1997.
- Terence Taylor. *Escaping the Prison of the Past: Rethinking Arms Control and Non-Proliferation Measures*. 1996.
- Terence Taylor and L. Celeste Johnson. *The Biotechnology Industry of the United States. A Census of Facilities*. 1995.
- Dean A. Wilkening. *The Evolution of Russia's Strategic Nuclear Forces*. 1998.
- Dean A. Wilkening. *How Much Ballistic Missile Defense Is Enough?* 1998.
- Dean A. Wilkening. *How Much Ballistic Missile Defense Is Too Much?* 1998.
- Dean A. Wilkening. *A Simple Model for Calculating Ballistic Missile Defense Effectiveness*. 1998.
- Zou Yunhua. *China and the CTBT Negotiations*. 1998.
- Zou Yunhua. *Chinese Perspectives on the South Asian Nuclear Tests*. January 1999.
- Chi Zhang et al. *Impact on Global Warming of Development and Structural Changes in the Electricity Sector of Guangdong Province, China*. 2000.

MacArthur Consortium Working Papers in Peace and Cooperation

- Pamela Ballinger. *Slaughter of the Innocents: Understanding Political Killing, Including Limited Terror but Especially Large-Scale Killing and Genocide*. 1998.
- Pamela Ballinger. *Claim-Making and Large-Scale Historical Processes in the Late Twentieth Century*. 1997.
- Tarak Barkawi. *Democracy, Foreign Forces, and War: The United States and the Cold War in the Third World*. 1996.
- Byron Bland. *Marching and Rising: The Rituals of Small Differences and Great Violence in Northern Ireland*. 1996.
- David Dessler. *Talking across Disciplines in the Study of Peace and Security: Epistemology and Pragmatics As Sources of Division in the Social Sciences*. 1996.
- Lynn Eden and Daniel Pollak. *Ethnopolitics and Conflict Resolution*. 1995.

Daniel T. Froats, *The Emergence and Selective Enforcement of International Minority-Rights Protections in Europe after the Cold War*. 1996.

Robert Hamerton-Kelly. *An Ethical Approach to the Question of Ethnic Minorities in Central Europe: The Hungarian Case*. 1997.

Bruce A. Magnusson. *Domestic Insecurity in New Democratic Regimes: Sources, Locations, and Institutional Solutions in Benin*. 1996.