Cybersecurity Economic Issues: Clearing the Path to Good Practice

Shari Lawrence Pfleeger and Rachel Rue, *RAND Corp.*

This framework helps project managers compare economic models to make effective decisions about cybersecurity investment and the trade-offs between investment and protection.

Software project managers have limited project resources. Requests for security improvements must compete with other requests, such as for new tools, more staff, or additional testing. Deciding how and whether to invest in cybersecurity protection requires knowing the answer to at least two questions: What is the likelihood of an attack, and what are its likely consequences? Security analysts understand a system's vulnerability to potential cyberattacks fairly well, but to date, research on the economic consequences of cyberattacks has been limited, dealing primarily with microanalyses of attacks' direct impacts on a particular organization. Many managers recognize the significant potential of a cyberattack's effects to cascade from one computer or business system to another, but there have been no significant efforts to develop a methodology to account for both direct and indirect costs. Without such a methodology, project managers and their organizations are hard pressed to make informed decisions about how much to invest in cybersecurity and how to ensure that security resources are used effectively.

In this article, we explore how others have sought answers to our two questions. We describe the data available to inform decisions about investing in cybersecurity and look at research models of the trade-offs between investment and protection. The framework we present can help project managers find appropriate models with credible data so that they can make effective security decisions.

Data realities

Many of us assume that a cyberattack's likelihood is reasonably high and might increase over time. Anecdotal evidence suggests that even organizations that have taken steps to detect and prevent attacks have experienced significant incidents. To provide a more realistic picture of the nature and number of cyber incidents, researchers have conducted several surveys in the last few years to capture information about security attacks and protection. The following are among the most well known:

- Since 2002, the annual Australian Computer Crime and Security Survey (ACC; www.auscert.org.au/render.html?it=2001) has used information provided by Australia's federal, state, and territorial law enforcement agencies and the national computer emergency response team AusCERT (www.auscert.org.au). It solicits data from large organizations about computer network attacks and computer misuse trends in Australia.
- The UK Department of Trade and Industry has administered several Information and Security Breaches Surveys, or ISBSs (www.infosec.co.

0740-7459/08/$25.00 © 2008 IEEE

uk/files/DTI_Survey_Report.pdf), since 1991. They report on Internet use, dependence on information technology, and computer security incidents at UK businesses.

- The annual CSI (formerly CSI/FBI) Computer Crime and Security Survey (http://i.cmpnet.com/v2.gocsi.com/pdf/CSISurvey2007.pdf) polls computer security practitioners in US corporations, government agencies, financial institutions, medical institutions, and universities that have joined the Computer Security Institute or attended a CSI seminar or workshop. The survey addresses computer usage, attacks, and actions taken in response to security incidents.
- Particular sectors have their own global surveys, such as the Deloitte-Touche Global Security Survey (www.dtti.com/dtt/article/0,1002,cid%253D172320,00.html). For example, the third GSS, administered in 2005, solicited input from chief security officers and security management teams of financial services industry organizations worldwide, asking for their perceptions of how one organization's information security compares to its counterparts' security.

Many surveys agree about which attacks are most serious: viruses, Trojan horses, worms, and malicious code.

These surveys paint a mixed picture. The ACC reported a decrease in attacks of all types at the same time that the ISBS found the percentage of attacked UK businesses had increased by a third over two years, and 43 percent of the CSI member organizations surveyed experienced increases in the rate of attacks. During the same period, the Deloitte-Touche survey found that the rate of financial-sector security breaches in the US had remained roughly the same. This variation in results might derive from the different populations being surveyed; the surveys represent different countries, sectors, degrees of sophistication about security matters, and biases in the pool of respondents. Moreover, most are convenience surveys (of self-selected respondents), so the population represented is unclear. This makes generalizing the results difficult.

Standard terminology

Another significant problem is the lack of standards in defining, tracking, and reporting security incidents and attacks. Surveys ask variously about the incidence of

- "electronic attacks" (ACC);
- "virus encounters" and "virus disasters" (ICSA Labs 8th, formerly the International Computery Security Association, Annual Computer Virus Prevalence Survey);
- "total number of electronic crimes or network, system, or data intrusions" (CSI/FBI);
- "security incidents," "accidental security incidents," "malicious security incidents," and "serious security incidents" (ISBS);
- "any form of security breach" (Deloitte);
- "unauthorized use of computer systems" (CSI/FBI); and
- "incidents that resulted in an unexpected or unscheduled outage of critical business systems" (Ernst and Young Global Information Security Survey, or EY).

It would be difficult to find two surveys with results that are strictly comparable. Thus, much of the reported evidence is categorized differently from one study to another, and the answers are based on respondents' opinions, interpretations, or perceptions, not on consistent capture and analysis of solid empirical data. This hodgepodge of definitions and concepts makes it difficult for software managers to know what data to collect and how to compare them with survey results.

The source and effects of attacks

Understanding the source of attacks is similarly problematic. For example, the ACC survey reports that the rate of insider attacks has remained constant. However, Deloitte claims that, in financial services, most attacks comes from the inside, and the rate is rising. The EY survey also emphasizes the rising threat of insider attacks. Several other surveys note that the sources of attacks are unknown in a significant percentage of cases. Surveys generally agree about which attacks are most serious: viruses, Trojan horses, worms, and malicious code. Most sectors also fear insider misuse and access abuse. Phishing is a growing concern, with general agreement that these attacks have increased dramatically over the past two years.

In addition to the number and kind of attacks, surveys often ask about effect, particularly in terms of cost. Significant variations exist in this category as well. For example, the ICSA survey has reported a 25 percent increase in the cost of recovering lost or damaged data. On the other hand, the ACC, EY, and CSI/FBI surveys found a decrease in total damage from attacks, even though this cost is increasing for some kinds of attacks (such as unauthorized access and theft, noted by CSI/FBI). Twenty-five percent of organizations reported financial loss to CSI/FBI, and 56 percent reported operational losses. Software managers need this cause-and-effect information, not only to design more secure systems but also to estimate resource needs for preventing, mitigating, and recovering from attacks, particularly attacks against the development platforms on which new software is being created.

Once again, the reasons for variations in findings are partly attributable to disparities in the respondent pools. However, a more significant problem, acknowledged both by survey respondents and administrators, is the difficulty in detecting and measuring both direct and indirect costs from security breaches. There are neither accepted definitions of loss nor standard, reliable methods to measure it. For example, the ICSA 2004 survey notes that "respondents in our survey historically underestimate costs by a factor of 7 to 10."

The various surveys do, in fact, reach consensus in several areas. For example, many surveys indicate that formal security policies and incident response plans are important. In addition, the lack of staff education and training within the IT security team and throughout the development organization appears to be a major obstacle to improved security. More generally, a poor "security culture" (in terms of awareness and understanding of security issues and policies) is often reported to be a problem and is a key concern of chief information security officers.[1] Survey respondents report that regular testing and updating of security procedures, combined with practices that increase staff awareness, are important. But little quantitative evidence supports these views, plausible as they may be.

The survey results also highlight another gap in our understanding of security investments: It's unclear how much organizations have invested in security protection, prevention, and mitigation. We know little about how they make investment decisions or measure their security investments' effectiveness; what little we do know is anecdotal and seems to be related to the business model and corporate culture.[2] Inputs required for such decision making—such as rates and severity of attacks, cost of damage and recovery throughout the enterprise, and actual cost of security measures of all types—are not known with any accuracy. Neither is it clear whether traditional measures, such as return on investment or internal rate of return, are appropriate for assessing security effectiveness. Simple questions, such as how much more security an extra dollar buys, go unanswered. To address some aspects of this situation, the US Bureau of Justice Statistics has administered the first large-scale, carefully designed and sampled cybersecurity survey. The results will provide the first official US statistics on the extent and consequences of cybercrime against US businesses.

Data needs

Software project managers need better data to support their decision making about security. Ideally, a data source should provide information to support the following tasks.

First, project managers must decide how to allocate resources to monitor and address cyber incidents. Survey data can inform resource allocation decisions ranging from monitoring cyber crime to sensing attempts at system penetration. Trend data about cyber incidents, including records from incident response teams, can support more effective strategic planning. Such data can then be used to help in

- understanding the project's current security practices;
- evaluating existing regulations and standards for incorporation into the system design;
- choosing effective measures of effectiveness for resource allocation;
- choosing organizational structures to facilitate efficient use of resources; and
- understanding current and future types and frequency of attack, probable and possible consequences for each type of attack, targeted sectors and businesses, and motives for and intended consequences of attacks.

Second, government, industry, and monitoring organizations must implement standards and guidelines. Trend data about vulnerabilities and attacks suggest areas that new or updated standards and guidelines might address. For example, the Common Vulnerabilities and Exposures list (http://cve.mitre.org) uses standardized names for vulnerabilities, which enables users to cross-reference and catalog them. The application of such standards lets disparate project management organizations search for common problems and possible solutions. In turn, standard classification and naming of vulnerabilities, types of attack, and techniques used in attacks can permit cross-project analysis that suggests best practices involving the most cost-effective technologies, policies and procedures, and organizational structures and processes.

Also, the insurance industry could play a growing role in securing cyberspace. For example, the Basel II agreement (standards originated by the Bank for International Settlements governing the amount of capital internationally active banks must have in reserve to guard against financial and operational risks) enables businesses to decrease their financial reserves in exchange for

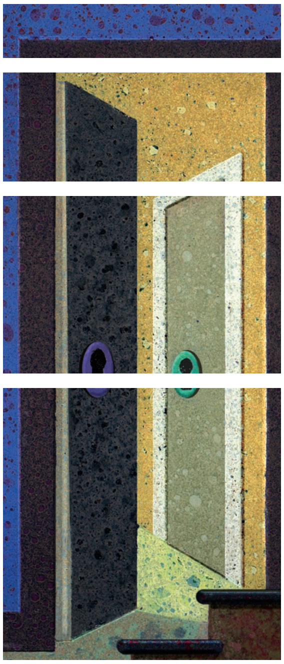

Anyone who has to choose and implement measures to increase security needs measures of effectiveness.

sharing information about cyber vulnerabilities and agreeing to comply with minimum standards. Credible survey data could be used to set policy terms and standards for insurability against cyberattacks. This information would inform decisions about how much security to build into a product and how much it would cost.

Fourth, there's a need for infrastructure protection benchmarks. Much critical infrastructure depends heavily on information technology. Repeated, coordinated surveys could be used to compare the cybersecurity postures of different infrastructures and to measure improvement. Benchmarks could support the analysis of attack frequency and severity trends and of consequent losses, determination of best practices for addressing current and changing vulnerabilities, and regular updating of standards.

Finally, anyone in the position of choosing and implementing measures to increase security needs measures of effectiveness. For example, survey data could provide feedback on the effectiveness of campaigns to strengthen cybersecurity. Data could also influence

- perception and empirical measurement of security strategies' effectiveness;
- development and dissemination of good metrics;
- perceived and actual effects of regulations and standards, and their enforcement; and
- perceived and actual effects of both public- and private-sector education strategies.

Research needs vs. realities

Multidisciplinary research is needed within and across the boundaries of engineering, business, and liberal arts and science. Indeed, social scientists have recently brought new insight to the traditional engineering problem of cybersecurity, with strong arguments for richer economic analysis. Ross Anderson notes that studying economic incentives is as important as studying the underlying technology—a call to technology-trained project managers to broaden their perspective.[3] But the literature is immature, with both a paucity of empirical analysis and a lack of agreement on findings. Researchers are pursuing the following four key streams of inquiry.

Software quality

In addition to technical deficiencies, business exigencies and market factors contribute to low software quality. Because software complexity can lead to faults that create vulnerability to attack, researchers have been examining the contextual factors, such as time to market and shareholder value, that drive quality investment decisions.

- *Patches.* Ashish Arora, Jonathan Caulkins, and Rahul Telang demonstrate that vendors have incentives to release products early and make repairs later using patches.[4] However, Dmitri Nizovtsev and Marie Thursby suggest that disclosing vulnerabilities and releasing fixes quickly tends to minimize total losses from attacks.[5]
- *Disclosure timing.* Questions remain about when and whether software testers and maintainers should publicly disclose newfound vulnerabilities. Researchers disagree about the details of specifying the vulnerabilities in a formal model.
- *Stock-price effects.* Corporations by law must create shareholder value, so researchers look at stock price for insight on vendor behavior, with mixed results. For example, Katherine Campbell and her colleagues found limited evidence of a negative market reaction to public announcements of security breaches.[6] By contrast, Rahul Telang and Sunil Wattal describe the drop in share price that results immediately after a software failure's disclosure.[7]
- *Vulnerability reduction.* Researchers also question the value of actively searching for vulnerabilities. Again, the results are mixed. Andy Ozment drew no clear conclusion from available data,[8] and Eric Rescorla found no relationship between software quality and the effort to remove vulnerabilities.[9]

Market interventions

Researchers are also examining the impact and effectiveness of information sharing, market mechanisms, and new approaches to insurance and liability.

Information-sharing programs. As with software quality, researchers are examining the economic consequences of and motivations for sharing security information. Bruce Schneier looks at full disclosure in the context of "inevitable vulnerability," proposing efforts to shrink the exposure window by prompting vendors to act quickly.[10] By drawing on literature from trade associations and research joint ventures, Lawrence A. Gordon and his colleagues illustrate a theoretical approach to analyzing the voluntary Information Sharing and Analysis Centers formed by American business sectors to advise the US government about homeland security issues.[11] Esther Gal-Or and Anindya Ghose go a step further, using a formal model of ISACs to show that security investments strategically complement information sharing.[12] These results provide a context for understanding when ISACs can be effective. Overall, researchers have found that information-sharing programs can have a significant and positive effect on security. Such findings suggest that software project managers should share security information whenever possible and practical.

Market mechanisms. Anderson draws parallels between cybersecurity and environmental pollution; both involve one group's investment benefitting another, a phenomenon called *negative externality*.[13] For this reason, researchers have proposed creating "vulnerability markets" using transferable security credits so that more vulnerabilities in one product can be balanced by fewer in other products. This concept affects how software products and product lines might be designed, with security trade-offs in vulnerability markets. For instance, Stuart Schechter proposes using vulnerability markets to benchmark security strength by rewarding bug discoveries.[14] He argues that such markets will yield information that can be used to improve the software development process. However, Ozment questions this proposal by comparing the mechanism to auctions, pointing out their inherent limitations.[15] Karthik Kannan and Telang also reject market-based mechanisms for two reasons: Their expected user losses would be higher than those of systems such as CERTs that rely on "benign identifiers" to report vulnerabilities and—even worse—they would provide incentives for misuse of vulnerability information by monopolists of an application or operating system.[16]

Liability reform. Some researchers recommend stricter liability requirements for software companies, ending the use of end-user license agreements that obviate developer responsibility. For example, Schneier suggests tightening liability on software manufacturers,[10] while Adam Shostack calls for more subtle use of vendor liability requirements to create better descriptions of product quality.[17] Jay Kesan, Ruperto Majuca, and William Yurcik explain the limitations of using traditional insurance to transfer cybersecurity risk, citing a lack of good data, overpricing, and excessive exclusions to skirt moral hazards.[18] Other researchers have identified additional challenges to a healthy insurance market, such as interdependent risk, which decreases the benefit of risk diversification. This condition results from widespread incentives to undersecure and the prevalence of dominant software packages, where a single exploitation can affect a large population of systems. Walter S. Baer and Andrew Parkinson explore other pros and cons of cyber insurance.[19]

Evaluation

It's not yet clear which evaluative criteria are most useful; no method has emerged as a gold standard. Initial calls to discard heuristics were replaced by insistence on *return on security investment* analysis. However, although consistent with other corporate investment decisions, ROSI and related concepts of internal rate of return and net present value are considered by some to be inappropriate frameworks for this kind of analysis. Lawrence Gordon and Martin Loeb contend that ROSI doesn't reveal the true economic rate of return and leads to the wrong investment objectives.[20] Huseyin Cavusoglu, Birendra Mishra, and Srinivasan Raghunathan suggest that ROSI is frustrated by the need to assign costs to poorly defined outcomes.[21]

Researchers are proposing new metrics to address cost assignment challenges. For example, Fariborz Farahmand and his colleagues consider using damage assessment across predefined categories as an evaluative framework.[22] Schechter introduces *cost-to-break* (that is, the effort required to invade a system) as a measure of security strength.[23] Cost-to-break and security strength work together to help model the effort an attacker must expend to gain access to a system. Schechter offers this measure to improve predictions about the amount of risk a system faces. Marco Cremonini and Patrizia Martini use the attacker's vantage point to evaluate the potential impact—and resulting benchmark for investment—of coupling a "Return on Security Investment" index with an estimate of the attacker's "Return On Attack."[24]

Enterprise decision making

Most information infrastructure is privately controlled, so economic researchers have provided tools—ranging from accepted methods instantiated in accessible tools to esoteric methods not yet widely available—to support analysis of corporate cybersecurity investments. For example, Gordon, Loeb, and Sohail Tashfeen provide a framework for thinking about trade-offs between investing indirectly in cyber insurance and directly in security countermeasures.[25] Furthermore, Gordon and Loeb offer a systematic way to incorporate qualitative information into investment analysis by prioritizing using the Analytical Hierarchy Process to elicit user preferences[26]; their book explains how to perform return-on-investment calculations.[27] Similarly, other researchers propose using Monte Carlo simulation to support midlevel managers in forecasting uncertainties in security investment decisions.[28] Kevin Soo Hoo[29] and Farahmand and his colleagues[30] suggest more expansive frame-

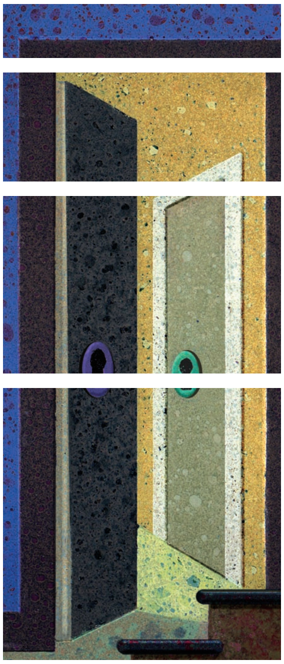

Table 1

List of characteristics used to describe cybersecurity economic models

Characteristic	Description
Type or form	The class of model and its mathematical structure
History and previous applications	When and for what purpose the model was originally developed and where it's been applied successfully
Underlying assumptions	Includes simplifications made to enable easier application
Decisions that the model supports	The types of decisions that a decision-maker would be able to substantiate by properly applying the model
Inputs and outputs	The quantities or attributes that the model manipulates
Parameters and variables	Elements that affect how the model transforms inputs to outputs
Applicable domain and range	Temporal and physical ranges of inputs, outputs, parameters, and variables that the model describes
Supporting data	Evidence that the model accurately represents the phenomena of interest

works using formal decision analysis to determine the amount to invest on the basis of the likelihood of intrusion. Cavusuglo, Mishra, and Raghunathan provide a detailed model to help companies select a security architecture on the basis of observed intrusions and derived likelihood of attack.[31]

A framework for comparing cybersecurity models

Thus, models abound, but project managers have trouble knowing which model to use and what credible data are available to inform it. To help managers understand each model's strengths and weaknesses and pick an appropriate model for each situation, we (along with David Ortiz) built the framework summarized in table 1.[32] The work of applying the framework to many classes of models is ongoing.

A model user must know whether the model's assumptions match the situation in which the model will be used. In general, the framework applies common accounting concepts, procedures, and principles, providing a platform for harmonizing different project-based modeling initiatives and data collection programs. It enables a manager to build a suite of models that address key cybersecurity investment decisions as each project is planned, implemented, and evaluated.

A key purpose of comparing models is to put them in a real-world context. Comparing models' strengths and weaknesses helps model developers and users understand the best ways to assemble needed data, run models, and present output and conclusions. You can use the following questions to contrast one model with another along several dimensions, each of which emphasizes the model's appropriateness for its intended use. The questions highlight the significance of model characteristics and help reveal gaps between models and the scenarios in which they're intended to be used.

Is the model relevant? Does it use data, methods, criteria, and assumptions appropriate for the reported information's intended use? Quantification of inputs and outputs should include only information that users (of the models and of the results) need for decision making. Data, methods, criteria, and assumptions that can mislead or that don't conform to carefully defined model requirements aren't relevant and shouldn't be included.

Is the model complete? Does it consider all relevant information that might affect the accounting and quantification of model inputs and outputs, and complete all requirements? Does it consider and assess all possible effects? Are all relevant technologies or practices considered as baseline candidates? The model's documentation should also specify how to collect all data relevant to quantifying model inputs.

Is the model consistent? Does it use data, methods, criteria, and assumptions that enable meaningful and valid comparisons? Methods and procedures should always be applied to a model and its components in the same manner. The same criteria and assumptions should be used to evaluate significance and relevance, and any data collected and reported should be compatible enough to allow meaningful comparisons over time.

Is the model transparent? Does it provide clear and sufficient information for reviewers to assess its credibility and reliability and the claims derived from it? Transparency is critical, particularly given the flexibility and policy relevance of many decisions based on the models' outputs. Information about the model and its usage should be compiled, analyzed, and documented clearly and coherently so that reviewers can evaluate its credibility. Specific exclusions or inclusions should be clearly identified, assumptions should be explained, and appropriate references should be provided for both data and assumptions. Information relating to the model's "system boundary" (the part of the problem addressed by the model), the identification of baseline candidates, and the estimation of baseline data values should be sufficient to enable reviewers to understand how all conclusions were reached. A transparent report will provide a clear understanding of all assessments supporting quantification and conclusions. This analysis should be supported by comprehensive documentation of any underlying evidence to confirm and substantiate the data, methods, criteria, and assumptions used.

Is the model accurate? Does it reduce uncertainties (with respect to measurements, estimates, or calculations) as much as is practical? Measurement and estimation methods should avoid bias. Acceptable levels of uncertainty will depend on the model's objectives and the intended use of results. Greater accuracy will generally ensure greater credibility for any model-based claim.

Is the model conservative? Does it use conservative assumptions, values, and procedures when uncertainty is high? Where data and assumptions are uncertain and where the cost of measures to reduce uncertainty isn't worth the increase in accuracy, conservative values and assumptions (those more likely to underestimate than overestimate changes from the baseline or initial situation) should be used.

Does the model provide insight? Does the model state clearly the nature of the insights it provides? Models might in some cases not serve to generate a specific result but rather to provide a means for decision-makers to understand the complex problems they face.

No single model can provide a comprehensive approach to guide cybersecurity investments. Indeed, because finding both credible data and an appropriate model is difficult, it's often unclear how a particular cybersecurity model can be used to support practical decision making. Software managers should understand how to evaluate and use several models in concert, either to triangulate and find an acceptable strategy for investing in cybersecurity or to address multiple aspects of a larger problem.

Cybersecurity economics is an emerging field. The first Workshop on the Economics of Information Security was held in 2002, and *IEEE Security and Privacy* devoted its first special issue to it in January 2005. All of us are eager for better data, better understanding, and better methods for using resources wisely in protecting critical products and services and providing assurance that software will work as expected. We can be active players in improving our understanding of cybersecurity economics by monitoring cyber incidents and responses, soliciting and using standard terminology and measures, and sharing data whenever possible. We can keep abreast of cybersecurity economics issues by participating in the Seventh Annual Workshop on the Economics of Information Security, to be held at Dartmouth College in 2008 (http://weis2008.econinfosec.org/index.htm). We can also participate in surveys and studies to better understand the nature and extent of cyber incidents. We can share information with researchers and colleagues to enable business sectors to take a coordinated approach to preventing and mitigating attacks. And finally, we can apply appropriate business measures to balance security investments with other requests for corporate resources.

Acknowledgments

This work was supported by the Economics of Cyber Security project of the Institute for Information Infrastructure Protection (www.thei3p.org) under award 2003-TK-TX-0003 from the US Office for Domestic Preparedness and Office of Justice Programs and the Department of Homeland Security. The presentation is based on RAND Corp. research and the authors' opinions.

References

1. M.E. Johnson and E. Goetz, "Embedding Information Security Into the Organization," *IEEE Security and Privacy*, vol. 5, no. 3, 2007, pp. 16–24.
2. S.L. Pfleeger, M. Libicki, and M. Webber, "I'll Buy That! Cybersecurity in the Internet Marketplace," *IEEE Security and Privacy*, vol. 5, no. 3, 2007, pp. 25–31.

About the Authors

Shari Lawrence Pfleeger is a senior information scientist at the RAND Corp., specializing in software engineering, cybersecurity economics and measurement, and IT policy. She coauthored, with various colleagues, *Security in Computing*; *Software Engineering: Theory and Practice*; *Solid Software*; *Applying Software Metrics*; and *Software Metrics: A Rigorous and Practical Approach*. She is a senior member of the IEEE Computer Society and the ACM and an editor for *IEEE Security and Privacy* magazine. She received her PhD in information technology and engineering from George Mason University. Contact her at RAND Corp., 1200 South Hayes St., Arlington, VA 22202-5050; pfleeger@rand.org.

Rachel Rue is an associate operations research analyst at RAND with research interests in cryptography, mathematical modeling, and information assurance. She is a member of the ACM and the Association for Unmanned Vehicle Systems International. She received her MS in algorithms, combinatorics, and optimization from Carnegie Mellon University and her PhD in philosophy from Princeton University. Contact her at RAND Corp., 4570 Fifth Ave., Pittsburgh, PA 15213; rue@rand.org; www.thei3p.org.

3. R. Anderson, “Why Information Security Is Hard—An Economic Perspective,” *Proc. 17th Ann. Computer Security Applications Conf.*, Assoc. for Economic Service, 2001, pp. 358–365.
4. A. Arora, J.P. Caulkins, and R. Telang, “Sell First, Fix Later: Impact of Patching on Software Quality,” Oct. 2004, http://ssrn.com/abstract=670285.
5. D. Nizovtsev and M. Thursby, “Economic Analysis of Incentives to Disclose Software Vulnerabilities,” paper presented at 4th Ann. Workshop Economics of Information Security (WEIS 05), 2005, www.infosecon.net/workshop/pdf/20.pdf.
6. K. Campbell et al., “The Economic Cost of Publicly Announced Information Security Breaches: Empirical Evidence from the Stock Market,” *J. Computer Security*, Mar. 2003, pp. 431–448.
7. R. Telang and S. Wattal, “Impact of Software Vulnerability Announcements on the Market Value of Software Vendors: An Empirical Investigation,” paper presented at 4th Ann. Workshop Economics of Information Security (WEIS 05), 2005, www.infosecon.net/workshop/pdf/telang_wattal.pdf.
8. A. Ozment, “The Likelihood of Vulnerability Rediscovery and the Social Utility of Vulnerability Hunting,” paper presented at 4th Ann. Workshop Economics of Information Security (WEIS 05), 2005, www.infosecon.net/workshop/pdf/10.pdf.
9. E. Rescorla, “Is Finding Security Holes a Good Idea?” paper presented at 3rd Ann. Workshop Economics of Information Security (WEIS 04), 2004, www.dtc.umn.edu/weis2004/rescorla.pdf.
10. B. Schneier, “Full Disclosure and the Window of Exposure,” *Crypto-gram Newsletter*, 15 Sept. 2000, www.schneier.com/crypto-gram-0009.html.
11. L.A. Gordon, M.P. Loeb, and W. Lucyshyn, “Sharing Information on Computer Systems: An Economic Analysis,” *J. Accounting and Public Policy*, vol. 22, no. 6, 2003, pp. 461–485.
12. E. Gal-Or and A. Ghose, “The Economic Incentives for Sharing Security Information,” *Information Systems Research*, vol. 16, no. 2, 2005, pp. 186–208.
13. R. Anderson, “Unsettling Parallels between Security and the Environment,” paper presented at the Workshop Economics of Information Security (WEIS), 2002, www2.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/econws/37.txt.
14. S. Schechter, “Computer Security Strength and Risk: A Quantitative Approach,” doctoral dissertation, Division of Eng. and Applied Science, Harvard Univ., 2004.
15. A. Ozment, “Bug Auctions: Vulnerability Markets Reconsidered,” paper presented at 3rd Ann. Workshop Economics of Information Security (WEIS 04), 2004, www.dtc.umn.edu/weis2004/ozment.pdf.
16. K. Kannan and R. Telang, “Market for Software Vulnerabilities? Think Again,” *Management Science,* vol. 51, no. 5, 2005, pp. 726–740.
17. A. Shostack, “Avoiding Liability: An Alternative Route to More Secure Products,” paper presented at 4th Ann. Workshop Economics of Information Security (WEIS 05), 2005, www.infosecon.net/workshop/pdf/44.pdf.
18. J. Kesan, R. Majuca, and W. Yurcik, “CyberInsurance as a Market-Based Solution to the Problem of Cybersecurity—A Case Study,” paper presented at 4th Ann. Workshop Economics of Information Security (WEIS 05), 2005, http://infosecon.net/workshop/pdf/42.pdf.
19. W.S. Baer and A. Parkinson, “Cyberinsurance in IT Security Management,” *IEEE Security and Privacy*, vol. 5, no. 3, 2007, pp. 50–56.
20. L. Gordon and M. Loeb, “Return on Information Security Investments: Myths versus Realities,” *Strategic Finance*, vol. 84, no. 5, 2002, pp. 26–31.
21. H. Cavusoglu, B. Mishra, and S. Raghunathan, “The Effect of Internet Security Breach Announcements on Market Value: Capital Market Reactions for Breached Firms and Internet Security Developers,” *Int’l J. Electronic Commerce*, vol. 9, no. 1, 2004, p. 70–104.
22. F. Farahmand et al., “A Management Perspective on Risk of Security Threats to Information Systems,” *Information Technology and Management*, vol. 6, nos. 2–3, 2005, pp. 203–225.
23. S. Schechter, “Quantitatively Differentiating System Security,” paper presented at 1st Workshop Economics of Information Security, 2002, www2.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/econws/31.pdf.
24. M. Cremonini and P. Martini, “Evaluating Information Security Investments from Attackers Perspective: The Return-on-Attack,” paper presented at 4th Ann. Workshop Economics of Information Security (WEIS 05), 2005; www.infosecon.net/workshop/pdf/23.pdf.
25. L. Gordon, M. Loeb, and S. Tashfeen, “A Framework for Using Insurance for Cyber-Risk Management,” *Comm. ACM*, vol. 46, no. 3, 2003, pp. 81–85.
26. L. Gordon and M. Loeb, “Evaluating Information Security Investments Using the Analytical Hierarchy Process,” *Comm. ACM*, vol. 48, no. 2, 2005, pp. 78–83.
27. L. Gordon and M. Loeb, *Managing Cybersecurity Resources*, McGraw-Hill, 2005.
28. J. Conrad, “Analyzing the Risks of Information Security Investments with Monte-Carlo Simulations,” paper presented at 4th Ann. Workshop Economics of Information Security (WEIS 05), 2005, http://infosecon.net/workshop/pdf/13.pdf.
29. K. Soo Hoo, “How Much Is Enough? A Risk-Management Approach to Computer Security,” Consortium for Research on Information Security and Policy, Stanford Univ., 2000, http://iis-db stanford.edu/pubs/11900/soohoo.pdf.
30. F. Farahmand et al., “Assessing Damages of Information Security Incidents and Selecting Control Measures: A Case Study Approach,” paper presented at 4th Ann. Workshop Economics of Information Security (WEIS 05), 2005, www.infosecon.net/workshop/pdf/39.pdf.
31. H. Cavusoglu, B. Mishra, and S. Raghunathan, “A Model for Evaluating: IT Security Investments,” *Comm. ACM*, vol. 47, no. 7, 2004, pp. 87–92.
32. R. Rue, S.L. Pfleeger, and D. Ortiz, “A Framework for Classifying and Comparing Models of Cyber Security Investment to Support Policy and Decision Making,” paper presented at 2007 Workshop Economics of Information Security (WEIS 07), 2007, http://weis07.infosecon.net/papers/76.pdf.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.