The State of Economics of Information Security

L. JEAN CAMP

ABSTRACT

This article discusses the emergence of a new area of study, known as economics of information security, by describing the initial work in this field. The article notes that economics of information security utilizes technical, business, policy, and applied perspectives, and workshops focused on economics of information security facilitated the coordination of work in this area of study. Next, the article provides an overview of a selection of the work in economics of information security's major areas of inquiry, which include: the role of insurance, the optimal construction of a market for vulnerabilities, the strategic role of security in the firm, the economics of privacy, the role of individual incentives, and the economics of digital rights management. Finally, this article introduces four of the current contributions to the field of economics of information security, which appear in this I/S issue on cybersecurity.

INTRODUCTION

This inaugural I/S symposium on cybersecurity policy helps to mark the increasing maturation of "Economics of Information Security" as an emerging area of study. The economics of information security, an explicit combination of primary disciplines, is *cross-disciplinary* as much as *interdisciplinary*. This overview is intended to provide a snapshot of the field as it stands and to identify a number of the critical questions likely to occupy researchers in the near term.

Economics of information security was initiated in a nearly simultaneous and completely uncoordinated manner at four institutions. In 2000, scientists at the Computer Emergency Response Team (CERT) at Carnegie Mellon proposed an early mechanism for risk assessment. The Hierarchical Holographic Model provided the first multifaceted evaluation tool to guide security investments using the science of risk. [21] Since that time, CERT has developed, under the name OCTAVE, a suite of systematic mechanisms for organizations to use in risk evaluations, depending on the size and expertise of the organization.

Shortly before this, Catherine Wolfram and I [8], from Harvard's Department of Economics and School of Government, respectively, had published an article employing careful economic definitions to define the specific "good" that is now widely considered the medium of exchange in the various theoretical constructions of security markets. Vulnerabilities were defined in this work as tradable externalities. Today, the market in vulnerabilities is quite real, with the purchase of a zero-day exploit by 3Com from an anonymous hacker in October 2005. [11] The normal count of days is from announcement of a vulnerability to exploit. A zero-day exploit is a vulnerability that has not yet been recorded as having been used in an attack. Microsoft claims to have discovered another zero-day exploit. [33]

In 2001, Ross Anderson of Cambridge published, "Why Information Security is Hard: An Economic Perspective," [2] at the Cambridge University Computer Laboratory. Professor Anderson explained that a significant difficulty in the optimal development of security technology is the imperative to integrate economic implications into technical designs. But, if a security technology requires that the party with the least risk make the greatest investment, then that system will fail to be widely adopted.

Also in 2001, Larry Gordon and Marty Loeb published a framework on "Using Information Security as a Response to Competitor Analysis Systems." [15] These professors at the University of Maryland's Smith School of Business examined the strategic use of security information from a classical business perspective.

A fifth notable work appeared in the business press, authored by the widely respected Dan Geer. [13] From his position as a recipient of privileged information on business investments at stake, he developed an argument for security investment to be measured not strictly by technical measures such as hardening or a simple count of dollars invested, but through a systematic Return on Security Investment Analysis.

These works together laid the foundation for an investigation of the economics of information security from technical, business, policy, and applied perspectives. The variety of schools and researchers engaged from these serendipitous beginnings has steadily expanded.

EMERGENCE OF ECONOMICS OF INFORMATION SECURITY

The disconnected but harmonious work published by 2001 indicated the potential of a new arena of intellectual endeavor, which might genuinely inform policy. Yet four articles do not make a body of knowledge. Bringing economics to bear upon the pressing questions of securing the commercial, academic, public, and personal networks that connect to form the national infrastructure required a more coordinated approach.

Ross Anderson and Hal Varian spearheaded the needed coordination by convening the Inaugural Workshop on the Economics of Information Security in at the University of California-Berkeley in 2002. These professors accompanied their invitations to the authors of all the previously mentioned work with an open call for papers. The inaugural workshop organized the discrete investigations into a set of core inquiries:

- The role of insurance;
- The construction of a market for vulnerabilities;
- The strategic role of security in the firm, including investments and disclosures;
- The economics of privacy as distinct from security;
- The individual role, as distinct from the national or firm; and
- The economics of digital rights management.

The major discovery of the first workshop was the variety of approaches and the wealth of current, but previously unorganized research. From Harvard, Stuart Schechter developed an innovative metric: the cost to break into a system. The cost to break, as opposed to classical risk analysis, provides a quantifiable measure of improvement in order to evaluate the Return on Security Investment Analysis. [13] From Maryland, Gordon illustrated that information sharing organizations are valuable, even in the case when some participants provide dishonest or incomplete information. [17] His focus was on the analysis of Information Sharing Analysis Centers (ISACs).

The contributions of the first, second, and third workshops were filtered and compiled into a single edited text, Economics of Information Security. [9] All the papers that were presented at the inaugural conference and its successors, plus future calls for papers and notices can be found at http://www.infosecon.net.

As of 2005, there is a single narrative that leads the reader through the questions, methods, and findings of the economics of information security by Gordon and Loeb. [16] The focus on the methodological exploration of security investment makes this text appropriate not only for a course but also for the individual seeking a guided introduction to the topic. With Gordon and Loeb as a primary text, and Camp and Lewis as a reference text, the economics of information security has reached the point where it is now a well defined academic foundation for coursework.

SELECTED FINDINGS

There has emerged a body of common findings that are now well understood. While there is continuing research, there is also a developing agreement with respect to the most cogent areas of investigation. Of course, the market for vulnerabilities has passed theory, moved through research, and is now clearly instantiated. What follows is an overview of the economics of information security work. This overview necessarily fails to include all significant works; otherwise this would become an annotated bibliography. However, the major areas of inquiry are included.

1. WHAT IS THE ROLE OF INSURANCE IN THE ECONOMICS OF INFORMATION SECURITY?

Insurance is a mechanism for enforcing contributions to a shared good. By requiring a minimal investment, insurance can address a situation where every party's risk is a function of the lowest investment, and thus there is a clear economic argument that insurance is appropriate for security mechanisms when the reliability and robustness of those mechanisms depends upon the weakest link. [31] Security mechanisms that exhibit this behavior include authentication systems based on shared information and denial of service attacks, where one firm can be attacked because of the existence of a network of subverted machines.

Insurance has now taken a significant role as an incentive for investment in security, with Lloyd's of London offering the first specific information security policy in 2003. Network security policies are also embedded in more traditional loss policies. Requirements for backup facilities and recovery plans as elements of disaster recovery policies enable organizations to better respond to electronic disasters. Counterpane Internet Security, for example, currently evaluates a commercial firm to provide metrics to determine if the firm is risk-seeking or has invested rationally in security. Before this practice became commonplace, the founder of Counterpane Internet Security, Bruce Schneier, presented a mechanism for developing such metrics [29] and illustrative cases where the lack of incentive for one firm had created costs for other firms in the same industry.

2. WHAT IS THE OPTIMAL CONSTRUCTION OF A MARKET FOR VULNERABILITIES?

The determination of vulnerabilities as a good was an important first intellectual foundation on which much has been built. However, in terms of research, much remains to be seen about how to construct a security market.

One mechanism for ensuring security is to develop formal price mechanisms to guide investments. Consider a software package. Initially, before a package is widely used and tested, there is a low bounty for vulnerabilities. There are ever-increasing bounty amounts. A small bounty, perhaps $10,000 for the first person to illustrate vulnerability, would be an opening offer. [28] As time passes and the system owner becomes more certain of security, the bounty can be increased. When a vulnerability is found, the bounty resets.

An extension that has not been previously considered is the adoption of per company bonds on privacy or security policies. For those nations that have strong privacy laws, there is an enforced commitment to their privacy policies at the risk of fines. An equivalent risk could be created by posting privacy bonds, whereby companies that handle data are forced to pay individuals, or corporate customers, when data are shared in violation of a previous commitment if confidentiality is lost.

An alternative mechanism is an auction that allows a person with knowledge of a vulnerability to announce its existence, while others indicate a willingness to pay. [23] The advantage of an auction is that it provides coordination for those willing to pay. Those who would gain the greatest value from investing in vulnerability disclosure (i.e., those with the lowest cost/benefit ratio) can set their willingness to pay. Bidding in this case could be organized as a multiple-good Dutch auction, where every party pays the price set by the first "losing" bidder, and the vulnerability is disclosed to those parties who pay. Alternatively, the auction could be a "reverse auction," which would provide the vulnerability to those parties who value the knowledge more than the threshold set by the discovering party. In either case, the party that identifies vulnerabilities would be paid at least as much as in any single purchaser case, and no company would pay more than the value of the vulnerability to the company.

Of course the value of an auction, in coordinating those at risk, requires the underlying coordination and information of the auction itself. Thus, what the market now sees are not auctions, but vendors. The purchasers of vulnerabilities are not producers of software, but the sellers of security services. Security vendors who pay for vulnerabilities have perverse incentives. A vendor who purchases vulnerabilities for its own subscribers or participants has no reason to maintain the confidentiality of that vulnerability. Once protected, the individuals who pay for the vulnerability have an incentive to leak information to illustrate the value of their service. [5]

A second more detailed analysis of the study of software vulnerabilities looks at the result of these perverse motivations of individuals and firms using repeated interactions (i.e., game theory). Formal disclosure of vulnerabilities, even those that are known in the community, increases their use. Thus, there is a possible argument that not spreading formal information about vulnerabilities may be best. White hats create a negative externality for black hats (i.e., they make the bad guys work harder). Currently, excluding the firm Tipping Point, there is only reputation capital for compensation for white hats who would expose vulnerabilities. White hats who sell vulnerabilities to a single vendor lose some reputation capital. Markets will increase the incentive to investigate but will also increase exposure. The optimal market would be one where there was a single purchaser who excludes no party from the information. This suggests direct governmental participation as a purchaser and distributor of vulnerability information, perhaps through an incident response team or ISAC. [19]

A study of a set of honeypots, including Linux and Windows, illustrates that formal disclosure of vulnerabilities, even those that are known in the community, increases their use. Formal disclosure increased the use of an informally known vulnerability by .26 attacks per day (e.g., one roughly every four days per machine) on average. Simultaneous publication and patching increases the observed attempts at subversion using the announced vulnerability by 0.02 attacks a day (or one attack every fifty days on one machine). The number of attacks per day is a per machine average, as the honeypot had multiple machines. [4]

3. WHAT IS THE STRATEGIC ROLE OF SECURITY IN THE FIRM?

A firm's investment security is obviously a function of its risk, defined as the product of the loss that would be created if there were a compromise and the probability of the compromise occurring. More detailed modeling [14] illustrates that the optimal investment depends very much upon the probability function, not simply the absolute probability. In fact, the shape of the probability function may result in investments ranging from nothing to nearly 40 per cent of the potential loss. This finding underlies the importance of collecting a range of comprehensive data about incidents and network activity, as enumerated by Shari Pfleeger. [24]

There are risks to investing in security to the extent that investment includes information sharing. The risk of possible losses in consumer trust and reputation discourages firms from sharing security information. Yet further research has verified that information sharing is both economically valuable and a complement to security investment. This research into information sharing has shown that information sharing is most valuable in highly competitive markets, because it counters downward pressures on pricing. [12]

There are also immediate costs to a firm that suffers a loss of information integrity. In addition to the long-term loss of reputation, a security incident is associated with immediate loss of value. A study of capital market valuation and announced incidents found that a firm loses more than two percent of its market value within two days of a publicized incident. Notice that this documented capitalization loss for firms with announced vulnerabilities yields a total loss that is greater than that reported by the annual FBI survey on cybercrime. This carefully calculated finding suggests that, far from security hysteria, there is still a widespread lack of security concern. [10]

In contrast, an examination of computer security from the perspective of insurance suggests that current practices may be reasonable. Either there is overinvestment, in which there are no incidents, or there is underinvestment, in which case there are incidents. Effectively, an insurance model suggests responding to the level of risk, implying that the current reactive practices are reasonable. If the risks are distributed as typical insurance risk (flood insurance in New Mexico? Hurricane insurance in Indiana?) then underinvestment before an incident is reasonable. However, this finding depends very much upon the relationship between past trends and current risk. [14] Lack of tornadoes in New York state indicates no tornadoes in New York state; never having been hacked may simply indicate a considerable run of luck. The distribution of risk is a critical unknown in cyber-insurance.

4. WHAT ARE THE ECONOMICS OF PRIVACY?

Why is it that the same individuals who express concerns about privacy will behave in a manner that systematically exposes their information? Economics offers a set of sometimes subtle answers.

First, the privacy market does not have adequate signals. At the most fundamental level, “protecting privacy” is a vague promise. For example, the privacy-enhancing technology market boom of the nineties included privacy protection that ranged from Zero Knowledge's provably secure and private email to Microsoft Passport's concentration of information in one (insecure) location. [7]

Even when privacy can be defined and specified, for example, through machine-readable P3P policies, a signaling problem remains. This signaling problem has been described in formal mathematical terms, and illustrates that the market for privacy cannot function without an external forcing function. The model of a market with fluctuating numbers of reliable privacy-respecting merchants will not necessarily reach an equilibrium where it is efficient for consumers to read privacy policies. As the cost of investigating the privacy policy changes, merchants that (dis)respect their own policies enter the market, the reliability of what is read varies, and there is no stable self-reinforcing equilibrium under which consumers should read privacy policies. Direct incentives are required to protect privacy. The market by itself will not reach an equilibrium where privacy policies are readable, read, and reliable as long as there are firms that can prevaricate about privacy. [32]

Beginning with an examination of the marketplace as a whole, not simply the digital marketplace, an argument can be made that there is a strong market for privacy. Products from simple window shades (with unarguably limited aesthetic appeal) to locking mailboxes thrive in the physical realm. Observing the physical and virtual markets for products providing unobservability, Shostack and Syverson conclude that, "when privacy is offered in a clear and comprehensible manner, it sells." [30] The argument is supported by the documentation of a range of sources of possible privacy products, from curtains to cryptography, which demonstrate the scale of these markets.

The understanding of privacy information as unreliable, and the market for privacy information as flawed, provides an important element to understanding user behavior. Individuals react in an understandable manner when information about privacy protection is ill-defined, untrustworthy, or even invisible. Signals in the privacy market are rejected when they are no more enlightening than the left turn blinker of a speeding octogenarian.

Alternatively, end user behavior can be categorized as simply discounting privacy risks. Individuals may share information, be aware of the risks, and simply discount those risks. Individual risk behaviors in other contexts are well documented and irrational. Privacy has none of the characteristics that generate horror, which would make the risks seem high, and the ubiquity of information sharing has made the risk too commonplace to create tension. [6] The calculus of computer security risk enabled by the CERT OCTAVE methodology is unquestionably beyond the limits of most computer users, and security is arguably a subset of the question of privacy.

Data compiled from privacy behaviors suggest that whatever the risks and whatever the reason, the risks of privacy are in fact discounted in consumer decision-making. In fact, individuals not only immediately discount privacy risk, but they increase their discount rate over time. [1] This is particularly interesting considering the rapid rate of increase in identity theft that suggests that risks increase over time.

5. WHAT IS THE ROLE OF INDIVIDUAL INCENTIVES?

The previous work assumes that privacy is good for individuals and good in some cases for firms. Yet, the information market is not always a zero-sum game in which gains from the consumer are offset by loses for the firm. Sharing information that is good for one party may not be in the interest of the other party. Privacy can be good for individuals or bad, e.g., when the information obtained by others is used to lower prices or to extend privileges. In particular, the opposite of privacy in the market is not necessarily information; the opposite of privacy is price discrimination. In markets where there is zero marginal cost (e.g., information markets) firms must be able to extract consumer surplus by price discrimination. This means that the firms cannot charge what they pay, at the margin, but must charge what the consumer is willing to pay. Data the consumer considers to be privacy violations may be necessary pricing data to the merchant. [22]

Experiments on the willingness of individuals to share data show that the farther someone is from the average, the more that person wants to protect their privacy. [18] That is, if a person's weight, salary, or age is close to the mean in a group, that person would not ask for much money for disclosure. But if a person's weight, salary, or age are far from the mean, then that person would demand more money for disclosure. This finding was based on experimental psychology. However, information theory predicates that the further data are from the mean, the more the data have the potential to reduce uncertainty. Therefore, the two sets of insights together argue that individuals, when empowered, rationally price information. Indeed, further empirical work [34] indicates that users are quite sensitive to the implications of further sharing of data and data type, so this sensitivity to the mean may not be generalizable.

Individual rejection of security information may itself be rational. When information security means ensuring that the end user has no place to hide his or her own information, or when security is implemented to exert detailed control over employees, individuals rightly seek to subvert the control. Security is often built with perverse incentives. Privacy and security are constructed to be opposites instead of complements in controlling information. Rejection of security is, in some cases, strictly rational. [26]

ECONOMICS OF DIGITAL RIGHTS MANAGEMENT

The most direct and obvious point of opposition between consumers and producers of computer security occurs in the implementation of Digital Rights Management (DRM). DRM implements business plans and strategies in information goods. Thus, the economics of DRM is a specialized arena of significant importance.

The initial study indicated the cynic's worst fear, which is that security as implemented in DRM is in opposition to security in terms of the owner and operator of the machine. DRM limits user options and competition, while not contributing to the security of machines. [3] Examples include tying batteries to phones and cartridges to printers. To the extent that security promotes survivability and the capacity to function in the face of attack, DRM is in opposition to security. By examining the return on complementary products, the action of the firm in implementing such (pseudo) security can be well understood.

DRM is used when legal remedies, based on protection of intellectual property to prevent unfair exploitation of innovation, are not available. The implementation of DRM in these cases does not support innovation, but rather only lock-in. Careful observation of the optimal investment in terms of social welfare identifies social and consumer costs. Limitations on reverse engineering that serve only to prevent competition are counted in economic terms as wasteful. [25]

Content holders have invested in DRM with the hope that such technology will force consumers to spend more and limit consumer sharing of information. Economic models illustrate that the true implications of DRM may not be all that the proponents hope for. A simple, clear examination of the cost of DRM indicates that the purpose is to increase friction in the market. Thus, providers of content, in order to prevent free sharing of content (also called piracy, depending on the speaker), increase friction in the purchase and use of goods. Yet, the option of free downloads remains, despite lawsuits and technologies. Observations of other markets, for example, in software, illustrate that the only way to compete with free availability is to increase service and reduce friction. Every expenditure in DRM that results in a reduced service or increased friction is an investment that will drive users to free, illegal but usable alternatives. [20] Examining DRM in the larger economic context, rather than focusing on the narrow potential of enforcing a particular license post-purchase, illustrates the risks to producers of DRM.

In fact, trusted computing is often considered the DRM Holy Grail. An economic analysis suggests that trusted computing arguably will help those who illegally upload information more than those who would prevent free information sharing. Current efforts against large-scale illegal copying on peer-to-peer networks depends on being able to prosecute by (in technical terms) violating the confidentiality of the users. Trusted computing would create an environment where peer-to-peer systems provide confidentiality to those who upload files, as well as integrity of content. Therefore, in a network characterized by trusted computing, users of peer-to-peer systems would be better off while those attempting to hold them legally accountable would be prevented from identifying those uploading files. [27]

In summary, the economics of DRM have illustrated that the incentives of DRM technology may be perverse, and thus, the results are not in the interest of those who support DRM.

CONCLUSIONS

Sony Corporation added DRM in its music compact discs allegedly to prevent illegal copying, implicitly validating observations about the economic waste in DRM. In fact, the copy protection software took the form of malicious software, a root kit that installed regardless of the user's selections in the installation dialogue. By virtue of installation at the most fundamental authentication level, the root, the toolkit has in many cases more authority than the CD listener. This DRM radically reduced the consumers' ability to secure their own machines, thus confirming the arguments that users are right to avoid some instances of security. The DRM also sent information back to the Sony Corporation advertising bureau, to enable price discrimination and targeting of advertising to consumers despite the stated privacy policy, thereby confirming both the relationship to privacy and price discrimination, and the near-zero value of stated privacy policies. The Sony DRM root kit is a disaster in security terms, but a disaster that was completely predictable and perfectly explicable in terms of economics of security. Those studying the economics of information security were less than shocked to discover the nakedness of that particular emperor.

Economics of information security has the potential to inform privacy and security-related initiatives, such as DRM, from policy and economics perspectives. Following its now-confirmed tradition of crossdisciplinary publication, economics of information security is unified as an intellectual endeavor by a series of workshops. The best papers from these workshops develop into journal special issues and texts, such as the special issue you now read (and may even hold in physical instantiation).

The work by Granick was first presented at the Fourth Workshop on the Economics of Information Security at Harvard University. Granick shows that the current legal construction of computer crime does not provide either clear incentives to invest in security or disincentives to commit computer crime. In computer crime, the cost to the victim of the crime is determined by the victim of the crime both before and after the incident. Companies that are ill-prepared even to the point of negligence can point to all their response costs, even those created by their own processes, as caused by an intrusion. For example, a company that fails to have even a trivial firewall can point to the post-incident purchase of a firewall as a cost of intrusion, as opposed to being held negligent to the point of creating an attractive hazard. Companies that overrespond to the point of paranoia can similarly run up costs and thus the putative harm of the crime. The law arguably protects at least those organizations that are the most prepared before an event and are the most professional in response. The punishments, as currently defined, may fit neither crime nor criminal. The incentives under the law are perverse, and the market cannot reverse those incentives.

The work by Rowe addresses IPv6 adoption in the United States, which has the greatest wealth of IP addresses, and can generously be described as glacial. The failure to adopt IPv6 is a refusal to invest, and a refusal to coordinate. After all, buying IPv6 increases the difficulty of using a domain as a platform to attack others, but it does little to prevent attacks. The failure of IPv6 adoption is a security failure that can only be understood in economic terms.

The work by Sand illustrates that the organization that integrates privacy with daily practice obtains the most value. A preliminary version of this work was presented as a work in progress at the Fourth Workshop on the Economics of Security. He identifies a dynamic of low investment and then expensive remediation that applies to privacy (as well as security). He identifies the loss of personal privacy as also a worst case for business as well as the individual, because a total loss of privacy for the person generates a loss of control over business assets by the organization. Finally, Sand provides two alternative frameworks and a set of concerns addressed by both of those frameworks in order to guide system designers in integrating privacy into process and technology.

The third fourth paper, a contribution led by Nathan Good, is the first examination of installation interfaces of software consisting of or containing spyware. This investigation builds upon the observations about signaling and information flow in the market for privacy and security. The fundamental question of the adequacy of informing users upon installation is addressed in a series of carefully designed usability tests. The finding is that mutual assent, given the state of law and computer interaction design, is not meaningfully achievable. However, the study did find that while individuals may not alter their behavior when notified of security and privacy risks, individuals nonetheless obtain a better emotional state postinstallation when provided such notification. The combination of incentive modeling, legal studies, and experimentation in this work both informs the reader, and illustrates that the study of economics of security has much to contribute.

In six years, the economics of information security has evolved from a disparate idea of disconnected scholars to a body of inquiry with a set of open questions, methods, and findings sufficiently examined as a means to begin to inform policy.

REFERENCES

[1] Acquisti, A., and J. Grossklags. 2004. Privacy attitudes and privacy behavior. In *Economics of Information Security*. Vol. 12 of *Advances in Information Security*, eds. Camp, L. J., and S. Lewis, 165–178. Boston: Kluwer Academic Publishers.

[2] Anderson, R. 2001. Why information security is hard: An economic perspective. In *ACSAC '01: Proceedings of the 17th Annual Computer Security Applications Conference*, 358. Los Alamitos, CA: IEEE Computer Society.

[3] ———. 2003. Cryptography and competition policy: Issues with 'trusted computing.' In *PODC '03: Proceedings of the twenty-second annual symposium on Principles of distributed computing*, pages 3–10, New York, NY, USA, 2003. ACM Press.

[4] Arora, A., R. Krishnan, A. Nandkumar, R. Telang, and Y. Yang. 2004. Impact of vulnerability disclosure and patch availability: An empirical analysis. Paper presented at the Third Workshop on the Economics of Information Security, Minneapolis, MN. http://www.dtc.umn.edu/weis2004/telang.pdf.

[5] Arora, A., R. Telang, and H. Xu. 2004. Optimal policy for software vulnerability disclosure. Paper presented at the Third Workshop on the Economics of Information Security, Minneapolis, MN. http://www.dtc.umn.edu/weis2004/xu.pdf.

[6] Camp, L. J. Forthcoming. Mental models of security. *IEEE Technology and Society*.

[7] Camp, L. J., and C. Osorio. 2003. Privacy-enhancing technologies for internet commerce. In vol. 2 of *Trust in the Network Economy* eds. Otto Petrovic, Michael Ksela, Markus Fallenbock, and Christian Kittl, 317-332. New York: Springer-Verlag/Wein.

[8] Camp, L. J., and C. Wolfram. 2000. Pricing security. In *Proceedings of the CERT Information Survivability Workshop*, 31–39. Boston: CERT.

[9] Camp, L. J., and S. Lewis, eds. 2004. Vol. 12 of *Advances in Information Security*. Boston: Kluwer Academic Publishers.

[10] Cavusoglu, H. 2004. Economics of it security management. In *Economics of Information Security*. Vol. 12 of *Advances in Information Security*, eds. Camp, L.J., and S. Lewis, 71-83. Boston: Kluwer Academic Publishers.

[11] Espiner, T. 2005. Symantec flaw found by TippingPoint bounty hunters. *ZDNET UK*, October 14. http://news.zdnet.co.uk/internet/security/0,39020375,39230317,00.htm.

[12] Galor, E., and A. Ghose. 2004. The economic consequences of sharing security information. In *Economics of Information Security*. Vol. 12 of *Advances in Information Security*, eds. Camp L.J., and S. Lewis, 95-105. Boston: Kluwer Academic Publishers.

[13] Geer, D. 2002. Making choices to show ROI. *Secure Business Quarterly* 1(2): 1–5.

[14] Gordon, L. A..,and M. Loeb. 2004. The economic of information security investment. In *Economics of Information Security*. Vol. 12 of *Advances in Information Security*, eds. Camp, L.J., and S. Lewis, 105-127. Boston: Kluwer Academic Publishers.

[15] Gordon, L. A.,and M. P. Loeb. 2001. Using information security as a response to competitor analysis systems. *Commun. ACM* 44(9): 70–75.

[16] Gordon, L. A., and M. P. Loeb. 2005. *Managing Cybersecurity Resources: A Cost-Benefit Analysis*. New York: McGrawHill.

[17] Gordon, L. A., M. P. Loeb, and W. Lucyshyn. 2002. An economics perspective on the sharing of information related to security breaches: Concepts and empirical evidence. Paper presented at the Workshop on the Economics of Information Security, Berkeley, CA. Available at http://www.cpppe.umd.edu/Bookstore/Documents/EconomicPerspective_05.17.02.pdf.

[18] Huberman, B. A., E. Adar, and L. R. Fine. 2005. Valuating privacy. Paper presented at the Fourth Annual Workshop on the Economics of Information Security, Cambridge, MA. http://infosecon.net/workshop/pdf/7.pdf.

[19] Kannan, K., and R. Telang. 2004. An economic analysis of market for software vulnerabilities. Paper presented at the Third Workshop on the Economics of Information Security, Minneapolis, MN. http://www.dtc.umn.edu/weis2004/kannan-telang.pdf.

[20] Lewis, S. 2004. How much is stronger DRM worth? In *Economics of Information Security*. Vol. 12 of *Advances in Information Security*, eds. L. Jean Camp and Stephen Lewis, 53-58. Boston: Kluwer Academic Publishers.

[21] Longstaff, T. A., R. Pethia, C. Chittister, and Y. Y. Haimes. 2001. Are we forgetting the risks of information technology? *Computer* 33(12): 43–52.

[22] Odlyzko, A. 2004. Privacy, economics and price discrimination on the internet. In *Economics of Information Security*. Vol. 12 of *Advances in Information Security*, eds. L. Jean Camp and Stephen Lewis, 187-212. Boston: Kluwer Academic Publishers.

[23] Ozment, A. 2004. Bug auctions: Vulnerability markets reconsidered. Paper presented at the Third Workshop on the Economics of Information Security, Minneapolis, MN. http://www.dtc.umn.edu/weis2004/ozment.pdf.

[24] Pfleeger, S. L., R. Rue, J. Horwitz, and A. Balakrishnan. Forthcoming. Investing in cyber security: The path to good practice. *The RAND Journal*.

[25] Samuelson, P., and S. Scotchmer. 2002. The law and economics of reverse engineering. *Yale Law Journal* 111(7): 1575–1663.

[26] Sandrini, M., and F. Cerbone. 2004. We want security but we hate it. In *Economics of Information Security*. Vol. 12 of *Advances in Information Security*, eds. L. Jean Camp and Stephen Lewis, 213-224. Boston: Kluwer Academic Publishers.

[27] Schechter, S. 2004. Toward econometric models of the security risk from remote attacks. Paper presented at the Third Workshop on the Economics of Information Security, Minneapolis, MN. http://www.dtc.umn.edu/weis2004/schechter.pdf.

[28] Schechter, S. E. 2004. Computer security strength and risk: A quantitative approach. Ph.D. thesis, Harvard University. http://www.eecs.harvard.edu/~stuart/papers/thesis.pdf.

[29] Schneier, B. 2002. We don't spend enough (on security). Comment at the Workshop on the Economics of Information Security, Berkeley, CA. http://www2.sims.berkeley.edu/resources/affiliates/workshops/econsecurity/econws/18.doc.

[30] Shostack, A., and P. Syverson. 2004. What price privacy? In *Economics of Information Security*. Vol. 12 of *Advances in Information Security*, eds. L. Jean Camp and Stephen Lewis, 129-142. Boston: Kluwer Academic Publishers.

[31] Varian, H. 2003. System reliability and free riding. In *Proceedings of the ICEC 2003*, N. Sadeh, ed., 355–366. New York: ACM Press.

[32] Greenstadt, V., T. Greenstadt, R. Greenstadt, and D. Molnar. 2004. Why we can't be bothered to read privacy policies. In *Economics of Information Security*. Vol. 12 of *Advances in Information Security*, eds. Camp, L.J., and S. Lewis, 143-154. Boston: Kluwer Academic Publishers.

[33] Wang, Y., D. Beck, X. Jiang, R. Roussev, C. Verbowski, S. Chen, and S. King. Automated web patrol with strider honeymonkeys:

Finding web sites that exploit browser vulnerabilities. Paper presented at the Internet Society Conference Proceedings for the Network and Distributed System Security Symposium (NDSS). http://www.isoc.org/isoc/conferences/ndss/06/proceedings/html/2006/papers/honeymonkeys.pdf.

[34] Wathieu, L., and A. Friedman. 2005. An empirical approach to understanding privacy valuation. Paper presented at the Fourth Annual Workshop on the Economics of Information Security, Cambridge, MA. http://infosecon.net/workshop/pdf/WathFried_WEIS05.pdf.