The State of Economics of Information Security

L Jean Camp

March 9, 2006

1 Introduction

The economics of information security is an emerging area of study. The economics of information security is *cross-disciplinary* as much as *interdisciplinary*. Economics of information security is the explicit combination of primary disciplines.

Economics of informations security has the potential to inform security from policy and economics perspectives. Following its now confirmed tradition of cross-disciplinary publication, economics of information security is unified as an intellectual endeavor by a series of workshops. The best from those workshops develops into journal special issues and texts. This works reviews the organization and findings of economics of information security.

2 Early Findings

Economics of information security was initiated in a simultaneous and completely uncoordinated manner at four institutions.

Also in 2000, the scientists as the Computer Emergency Response Team at Carnegie Mellon proposed an early mechanism for risk assessment. The Hierarchical Holographic Model provided the first multi-faceted evaluation tool to guide security investments using the science of risk. [22] Since that time, CERT has developed a suite of systematic mechanism for organizations to use in risk evaluations, depending on the size and expertise of the organization under the name OCTAVE.

Shortly before this publications, in 2000, work was published by myself and Catherine Wolfram, [8], from Harvard's School of Government and the Department of Economics, respectively. In this publication, using careful economic definitions, we defined the "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. [10].

In 2001, Ross Anderson of Cambridge published, "Why Information Security is Hard - An Economic Perspective", [3] at the University of Cambridge Computer Laboratory. Professor Anderson explained that a significant difficulty in optimal development of security technology is that economic implications should be integrated into technical designs. If a security technology requires that the party with the least risk makes the greatest investment, that system will fail to be widely adopted.

Also in 2001, Larry Gordon and Marty Leob published "A framework on using information security as a response to competitor analysis systems" [16]. These professor 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 [12]. 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 by hardening or investment, but through a Return on Security Investment Analysis.

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

3 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 does 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 Berkeley in 2002. The professors invited the authors of all the previously-mentioned work, and issued an open call. The Inaugural workshop jelled the investigations into a set of core queries.

- 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
- Economics of Digital Rights Management

The major finding of the first workshop was the variety of approaches and the wealth of current, but previously unorganized research. From Harvard, Stuart Schecter 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. [12]. From Maryland, Gordon illustrated that information sharing organizations are valuable, even in the case when some participants provide dishonest or incomplete information. [14] His focus was on the analysis of ISACs.

The contributions of the first, second and third workshops were filtered and compiled into a single edited text, Economics of Information Security. [6] All the papers that were presented at the four events, 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 economic of information security by Gordon and Loeb. [17] 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.

Certainly wIth Gordon and Loeb as a primary text, and Camp and Lewis as a reference text, economics of information security has reached the point where is now a well defined academic foundation for coursework.

4 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 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.

What is the role of insurance in the economics of information security?

Insurance is a mechanism for enforcing contribution 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. [35] Security mechanisms which 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 since then has taken a significant role as an incentive for investment in security; with Lloyds of London offering the first specific information security policy in 2003. Network security policies are also embedded in more traditional loss policie. For example, requirements back-up 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 commercial firms to provide metrics to determine if the firm is risk seeking or has invested rationally in security. Before this practice began more common-place, the founder of Counterpane Internet Security, Schneier, presented a mechanism for developing such metrics [32] and presented cases where the lack of incentive for one firm had created costs for firms in the same industry.

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 about how to construct a security market.

On mechanism for ensuring security it to develop formal price mechanisms to guide investments. Consider a software packeges. Initially, before it is widely used and tested, there is a low bounty for vulnerabilities. There there are ever-increasing bounty amounts. A small bounty, for example $10,000 for the first person to illustrate vulnerability would start. [31] As time passed 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 which have strong privacy laws there is an enforced commitment to their privacy policies or the risk of fines. An equivalent risk could be created by posting privacy bonds; whereby companies that handle data are forced to pay individuals when data are shared in violation of a previous commitment, or corporate customers if confidentiality are lost.

One alternative mechanism is an auctions so that a person with knowledge of a vulnerability announces its existence, and then others indicate a willingness to pay. [24] the advantage oa an auction is the it provide coordination for those willing to pay. Those with the greatest value in investing in vulnerability disclosure (e.g., those with the lowest cost benefit ratio)

can set their willingness to pay. Auctions 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 thresh hold set by the discovering party. In either case the party that identifies vulnerabilities at least as much as any single purchaser, and no company would pay more than the value of the vulnerability to the company.

Of course the value of auction in coordinating those at risk requires the underlying coordination and information of the auction itself. Thus vendors as opposed to auctions is what has arisen. Instead of producers of software, the purchasers of vulnerabilities are the producers of security services. Security vendors who pay for vulnerabilities have perverse incentives. A vendor that purchases vulnerabilities for its own subscribers or participants has no reason to maintain the confidentiality of that vulnerability. Once protected; the individuals that pay for the vulnerability have an incentive to leak information to illustrate the value of their service. [37]

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, e.g., 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 [1] create a negative externality on black hats; i.e. they make the bad guys work harder. Currently, excluding 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 investigation but will also increase exposure. The optimal market would be one where there was a single purchaser that excludes no party from the information. This suggest direct governmental participation as a purchaser and distributor of vulnerability information, perhaps through an incident response team or ISAC. [34]

Experimental evidence indicates that formal vulnerability disclosure does indeed cause an increase in its use. 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. [2] Formal disclosure increased the use of a informally known vulnerabilities by .26 a day, on average. Simultaneous publication and patching increases observed attempts at subversion with the announced vulnerability by 02. attacks a day. The number of attacks per day is a per-machine average, as the honeypot had multiple machines.

[1] "White hats" are those who seek to protect machines and networks. "Black hats" seek to subvert machines and networks. "Grey hats" are those who are engaged either in a mix of activities or in activities that are disputed; for example, announcing vulnerabilities before there are patches in order to force quick vendor response.

[2] A honey pot is a machine that has been given random content with attractive filenames, for example, "creditCardNumbers". Attacks on these honeypots are then studied.

What is the strategic role of security in the firm?

The investment of a firm in security is obviously a function of the risk, defined by quotient of the loss that would be created if there were a compromise and the probability of the compromise occurring. More detailed modeling [15] 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% of the potential loss. This finding underlies the importance of collecting a range of comprehensive data about of incidents and network activity, as enumerate by Pfleeger [25].

There are risks to investing in security to the extent that investment includes information sharing. Possible loss of consumer trust and reputation risks discourage 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. [11]

There are also immediate cost 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 2% 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 lack of concern. [9]

In contrast, an examination of computer security from the perspective of insurance suggests that current practices may be reasonable. Either there is over-investment, in which there are no incidents, or there is under investment, 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. [2]

What are Economics of Privacy

Why is it the case 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. [3] At the most fundamental level, "protecting privacy" is a vague promise. For example, the privacy-enhancing technology

[3]Signals referes here to economic signals not electromagnetic signals. In economics, signal are difficult to falsify data that differentiates types when there are superficially similar parties who are in fact different types.

market boom of the nineties included privacy protection that ranged from Zero Knowledge's provably secure and private email to Microsoft's Passports 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 can not function without an external forcing fucntion. A model of the 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, and the reliability of what is read varies, 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 a equilibrium where privacy policies are readable, read and reliable as long as there are firms that can prevaricate about privacy. [36].

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, the authors conclude that, "when privacy is offered in a clear and comprehensible manner, it sells". [33]. The argument is supported by the documentation of a range of sources of possibly privacy products, from curtains to cryptography, and document 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 risk. Individual risk behaviors in other contexts are well documented, and irrational. Privacy has none of the characteristics that generate horror, thus making risks seem high, and the ubiquity of information sharing has made the risk to commonplace to create tension. [5] 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 suggests that whatever the risks and why ever 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 the risks increase over time.

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 sometimes a zero sum game; that is, 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 or bad for individuals, if the information obtained by others' is used to lower prices or to extend priviledges. 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. What are privacy violations to the consumer may be necessary pricing data to the merchant. [23]

Experiments on individual's willingness to share data shows that the farther someone is from the average, the more that person wants to protect their privacy. [20]. 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 together argue that individuals, when empowered, rationally price information. Indeed further empirical work [19] indicates that users are quite sensitive to the implications of further sharing of data.

Indeed, 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. [29]

Economics of Digital Rights Management

The most direct and obvious point of opposition between consumer and producer of computer security occurs in the implementation of DRM. Digital Rights Management implements business plans and strategies in information goods. Thus the economics of DRM is a specialized arena.

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. [4] 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 (psuedo) 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. Thus that 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 [27].

Content holders have invested in DRM with the hopes 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 remain despite lawsuits and technologies. Observations of other markets, for example software, illustrate that the only way to compete with free 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. [21]. 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. [4] An economic analysis suggest 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 system provide confidentiality to those who upload files, and 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. [30]

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

[4] Trusted computing places a hardware cryptographic base that enables content providers to assert that specific distributions of content be accessed only with specific DRM-embedding applications. Therefore, for example, no one could play a movie on a player that does not embed region codes.

5 Recent Findings

The recent Sony DRM effort, described in the next paragraph, begs for economic examination. Those studying the economics of information security were less than shocked to discover the nakedness of that particular emperor.

Sony Corporation added DRM in its music compact discs allegedly to prevent illegal copying. (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 users' selections in the 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 machine. (Thus confirming the arguments that users are right to avoid some instances of security.) The DRM also sent information back to the Sony Corp 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.

The work by Granick [18] at the Fourth Workshop on the Economics of Information Security at Harvard University predicts continued unintended consequences of the current (broken) legal structure. She illustrates that the current construction of computer crime does not provide 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 which 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 failed to have even a trivial firewall can point to the post-incident purchase of a firewall as a cost of an intrusion, as opposed to being held negligent to the point of creating an attractive hazard. Companies that over-respond to the point of paranoia can similarly run up costs and thus the putative harm of the crime. The law arguably protects least those organizations that are the prepared before an event, and 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 Sand [28] illustrates that the organization that integrates privacy with daily practice obtains the most value. This work was also 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 worse case for business as well as individual, because total loss of privacy for the person is 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.

IPv6 adoption in the United States, which has the greatest wealth of IP addresses, 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 increase the difficulty of using a domain as a platform to attack others, but do little to prevent attacks. [26] The failure of IPv6 adoption is a security failure that can only be understood in economic terms.

Nathan Good [13] recently offered 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 obtained a better emotional state post-installation 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 come from a disparate idea of disconnected scholars to a body of inquiry with a set of open questions, methods, and findings sufficiently examined as to begin to inform policy.

References

[1] Alessandro Acquisti and Jens Grossklags. Privacy attitudes and privacy behavior. In L Jean Camp and Stephen Lewis, editors, *Economics of Information Security*, volume 12 of *Advances in Information Security*, pages 165–178, New York, NY, 2004. Springer.

[2] Roger Adkins. An insurance style model for determining the appropriate investment level. In *Third Workshop on the Economics of Information Security*, Minneapolis, MN, USA, June 2004.

[3] R. Anderson. Why information security is hard-an economic perspective. In *ACSAC '01: Proceedings of the 17th Annual Computer Security Applications Conference*, page 358, Washington, DC, USA, 2001. IEEE Computer Society.

[4] Ross Anderson. 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.

[5] L Jean Camp. Mental models of security. *IEEE Technology and Society*, 2006.

[6] L Jean Camp and S Lewis. *Economics of Information Security*, volume 12 of *Advances in Information Security*. Springer, New York, NY, 2004.

[7] L. Jean Camp and Carlos Osorio. Privacy enhancing technologies for internet commerce. In *Trust in the Network Economy*, Berlin, DE, 2003. Springer-Verlag.

[8] L Jean Camp and Catherine Wolfram. Pricing security. In *Proceedings of the CERT Information Survivability Workshop*, pages 31–39, Boston, MA, USA, October 2000. CERT.

[9] Huseyin Cavusoglu. Economics of it security management. In L Jean Camp and Stephen Lewis, editors, *Economics of Information Security*, volume 12 of *Advances in Information Security*, pages 71–83, New York, NY, 2004. Springer.

[10] Tom Espiner. Symantec flaw found by tippingpoint bounty hunters. *ZDNET*, October 2005.

[11] Esther Gal-Or and Anindya Ghose. The economic consequences of sharing security information. In L Jean Camp and Stephen Lewis, editors, *Economics of Information Security*, volume 12 of *Advances in Information Security*, chapter 8, pages 95–105. Springer, New York, NY, 2004.

[12] Daniel Geer. Making choices to show roi. *Secure Business Quarterly*, 1(2):1–5, October 2005.

[13] Nathaniel Good, Jens Grossklags, David Thaw, Aaron Perzanowski, Deirdre Mulligan, and Joseph Konstan. User choices and regret: Understanding users decision process about consensually acquired spyware. *I/S A Journal of Law and Policy for the Information Society*, 2006.

[14] Lawrence A. Gordon. An economics perspective on the sharing of information related to security breaches: Concepts and empirical evidence. In *Workshop on the Economics of Information Security*, Berkeley, CA, USA, May 2002.

[15] Lawrence A. Gordon and Martin Leob. The economic of information security investment. In L Jean Camp and Stephen Lewis, editors, *Economics of Information Security*, volume 12 of *Advances in Information Security*, chapter 9, pages 106–123. Springer, New York, NY, 2004.

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

[17] Lawrence A. Gordon and Martin P. Loeb. *Managing Cybersecurity Resources: A Cost-Benefit Analysis*. McGraw-Hill, NY, NY, 2005.

[18] Jennifer Granick. Faking it: Criminal sanctions and the cost of computer intrusions. *I/S A Journal of Law and Policy for the Information Society*, 2006.

[19] Bernardo A. Huberman, Eytan Adar, and Leslie R. Fine. Am empirical approach to understanding privacy valuation. In *Workshop on the Economics of Information Security*, Cambridge, MA, USA, June 2005.

[20] Bernardo A. Huberman, Eytan Adar, and Leslie R. Fine. Valuating privacy. In *Workshop on the Economics of Information Security*, Cambridge, MA, USA, June 2005.

[21] Stephen Lewis. How much is stronger drm worth? In L Jean Camp and Stephen Lewis, editors, *Economics of Information Security*, volume 12 of *Advances in Information Security*, chapter 4, pages 53–58. Springer, New York, NY, 2004.

[22] Thomas A Longstaff, Rich Pethia, C Chittister, and Y Y Haimes. Are we forgetting the risks of information technology. *IEEE Computer*, pages 43–52, 2001.

[23] Andrew Odlyzko. Privacy, economics and price discrimination on the internet. In L Jean Camp and Stephen Lewis, editors, *Economics of Information Security*, volume 12 of *Advances in Information Security*, pages 187–212, New York, NY, 2004. Springer.

[24] Andy Ozment. Bug auctions: Vulnerability markets reconsidered. In *Third Workshop on the Economics of Information Security*, Minneapolis, MN, USA, June 2004.

[25] Shari Lawrence Pfleeger, Rachel Rue, Jay Horwitz, and Aruna Balakrishnan. Investing in cyber security: The path to good practice. *The RAND Journal*, 2006.

[26] Brent Rowe and Michael Gallaher. Could ipv6 improve network security? if so, at what cost? *I/S A Journal of Law and Policy for the Information Society*, 2006.

[27] Pam Samuelson and Suzanne Scotchmere. The law and economics of reverse engineering. *Yale Lw Journal*, pages 1575–1663, 2002.

[28] Peter Sand. The privacy value. *I/S A Journal of Law and Policy for the Information Society*, 2006.

[29] Mauro Sandrini and Ferdinando Cerbine. We want security but we hate it. In L Jean Camp and Stephen Lewis, editors, *Economics of Information Security*, volume 12 of *Advances in Information Security*, pages 213–224, New York, NY, 2004. Springer.

[30] Stuart Schechter. Towards econometric models of software security risks from remote attacks. In *Third Workshop on the Economics of Information Security*, Minneapolis, MN, USA, June 2004.

[31] Stuart E. Schechter. Computer security strength and risk: A quantitative approach. In *Workshop on the Economics of Information Security*, Berkeley, CA, USA, May 2002.

[32] Bruce Schneier. We don't spend enough (on security). In *Workshop on the Economics of Information Security*, Berkeley, CA, USA, May 2002.

[33] Adam Shostack and Paul Sylverson. What price privacy? In L Jean Camp and Stephen Lewis, editors, *Economics of Information Security*, volume 12 of *Advances in Information Security*, pages 129–142, New York, NY, 2004. Springer.

[34] Rahul Telang and Karthik Kannan. An economic analysis of market for software vulnerabilities. In *Third Workshop on the Economics of Information Security*, Minneapolis, MN, USA, June 2004.

[35] Hal Varian. System reliability and free riding. In N. Sadeh, editor, *Proceedings of the ICEC 2003*, pages 355–366, New York, NY, USA, 2003. ACM Press.

[36] Tony Vila, Rachel Greenstadt, and David Molnar. Why we cannot be bothered to read privacy policies. In L Jean Camp and Stephen Lewis, editors, *Economics of Information Security*, volume 12 of *Advances in Information Security*, pages 143–154, New York, NY, 2004. Springer.

[37] Hao Xu. Optimal policy for software vulnerability disclosure. In *Third Workshop on the Economics of Information Security*, Minneapolis, MN, USA, June 2004.