Workshop on Information Security Economics*

Date of workshop: January 18–19, 2007

Workshop Organizers:

Jean Camp, Indiana University
ljean@ljean.com
Alessandro Acquisti, Carnegie Mellon University
acquisti@andrew.cmu.edu

Workshop Sponsors:

The Institute for Information Infrastructure Protection
Dartmouth University

DIMACS Center for Mathematics and Theoretical Computer Science

Report Author:
Tyler Moore
University of Cambridge
Tyler.Moore@cl.cam.ac.uk

Date of Report: February 6, 2007

1 Workshop Goals and Design

What is the role of economics in computer security? What is the role of computer security and privacy in the market? Are there economic models that can better inform system design? How can research in the economics of security inform legal practice? How can the design of secure systems be improved by an understanding of legal allocations of liability? Is there a set of questions that must be answered for this cross-discipline to move forward? What sets of methods are best suited to those questions?

This DIMACS Workshop addressed these questions and worked toward a concrete research agenda for the exploding area of study that combines security, privacy, and economics. The workshop had two primary goals. The first goal was to enlarge the interest in the economics of information security by bringing together researchers already engaged in the field with scientists and investigators in other disciplines. We summarize research presented in Section 2. The second goal was to produce – through two breakout sessions – an inclusive, integrated research agenda for this field of study. We summarize the results of these breakout sessions in Section 3.

2 Summary of Presentations

2.1 Foundational Concepts

Some speakers discussed fundamental tools and concepts important to information security economics. In "Perspectives from Microeconomic Theory and Game Theory", Beth Allen surveyed different economic tools that might be helpful in modeling information security problems. She argued that Pareto optimality may be impractical goal for engineered systems and that tools from game theory can usefully be applied. Standard pricing mechanisms run into problems because of the proliferation of asymmetric information, perverse incentives and free-rider problems. Game theory, by contrast, can accommodate strategic behavior and interdependent decision making, both of which are rife in the provision of information security mechanisms.

*DIMACS was founded as a National Science Foundation Science and Technology Center. It is a joint project of Rutgers University, Princeton University, AT&T Labs-Research, Bell Labs, NEC Laboratories America, and Telcordia Technologies, with affiliated partners Avaya Labs, IBM Research,Microsoft Research, and HP Labs. This workshop also acknowledges support from the Institute for Information Infrastructure Protection (I3P). The I3P is managed by Dartmouth College, and supported under Award number 2003-TK-TX-0003 from the U.S. DHS, Science and Technology Directorate.

It remains an open question whether games should be modeled as cooperative or non-cooperative; furthermore, when the structure of interactions matter, network games may be required.

In "Incentive-Centered Design for Information Security", Rick Wash highlighted the strong role humans have in achieving security, from choosing which system to use to evaluating security information to actually using systems once deployed. He noted that most computer systems are designed with little regard to human behavior; systems designers resign themselves to accommodating 'unpredictable' user actions. Such resignation is wrong, argues Wash. Humans *do* respond to incentives in strategic settings. As an example, he examines various "proof-of-work" schemes to combat spam e-mail.

He presents a classic screening model, where the system cannot distinguish good email from bad. Proof-of-work schemes ask users to perform a screening test; the aim is for the test to be easy to do a limited number of times but onerous in large amounts. The test must satisfy four properties: (i) cost increases in task intensity, (ii) cost is supra-linear, (iii) cost is greater for bad users than good ones, and (iv) incremental cost of harder tasks is greater for bad users than good ones. While the original proposals may have satisfied the properties, Laurie and Clayton argued that botnets undermine properties (iii) and (iv). He then notes that a reputation-based framework attributed to Liu and Camp might work. In the end, it is unclear whether the effect on users is to stop spam or simply increase the burden on good users. Wash concludes by arguing that user incentives should be considered in designing systems to "keep bad stuff out" (e.g., spam, spyware) and "get good stuff in" (e.g., privacy-enhancing technologies).

Bruce Schneier argued that "The Psychology of Security" is important for better understanding how people make (typically poor) risk-management decisions. Security is a trade-off, and while people have intuitions about these trade-offs, they are often wrong. For instance, we overestimate spectacular or rare threats while underestimating common risks and threats that are slow to evolve over time. Schneier discusses experiments from prospect theory, for example where people are risk-averse when presented with opportunities to gain and risk-seeking when minimizing loss. In sum, while Rick Wash argued that it is beneficial to design systems where we consider users as rational and responding to strategic behavior, Schneier noted that we must also deal with peculiar irrationalities users display when managing security risks.

2.2 Improving Transparency in Information Security

In "Notice of Security Breaches as a Lightweight Regulation", Deirdre Mulligan examined the role of disclosure in promoting better security practices and higher investment by companies. She first discussed the Emergency Planning and Community Right to Know Act (EPCRA), which required companies to disclose the release of toxic chemicals. This legislation prompted a significant reduction in the amount of chemicals disclosed by changing the behavior of companies. She noted that one reason for its success is that, unlike baseline standards, the disclosure prompted a "race to the top" by companies. Mulligan then recounted the adoption of California bill AB 700, which mandated that companies notify consumers of security breaches involving personal data. She noted that the breach law, while not measuring security, did introduce a measure of failure that has enabled companies to quantify the price of security failure. However, Mulligan notes that there are important differences between AB 700 and its inspiration, EPCRA, that could limit the security breach law's success. For EPCRA, the data were highly structured, widely available (stored in an EPA-administered database) and triggered community involvement. By contrast, the privacy breach data are non-standardized and distributed, while disclosures by third parties have limited the ability of consumers to "vote with their feet". Finally, Mulligan posed an open question how other disclosure mandates beyond those related to privacy violations could be usefully applied to strengthening security investment.

In "Linking the Economics of Cyber Security and Corporate Reputation", Barry Horowitz quantifies the effect of breach disclosure laws like AB 700 on security investment. He argues that the impact varies depending on the extent to which a privacy breach harms a business's reputation. Banks, for instance, are more concerned with protecting its image through tougher information security measures than manufacturers are. He constructs a model to infer relative importance of information security investment for different industries. He uses data from breach disclosures identified through the news media and publicly-available revenue reports. The model finds that the finance industry invests six times as much as the retail sector and three times as much as the manufacturing sector. One novelty in this work is how it overcame a lack of data in the amount of security investments by reverse-engineering a model to use breach disclosures as input.

In "Information Security and IT Risk Management in the Real World: Results from Field Studies", Scott Dynes underscored the need for transparency to trigger security investment by companies across many sectors.

Dynes described an information-gathering methodology where he investigated a host firm along with its suppliers to identify supply-chain dependencies on the Internet infrastructure. He notes that firms primarily take a local view to information security, not considering the effects on other companies in its sector or hidden liabilities that may exist at suppliers. He argues that while latent market forces do exist, not enough information about threats is being disclosed which hinders rational security investment.

In "Data Policy Violations", Dan Geer identifies data security as an important, but neglected, focus of study. He points to the increasing value of data as well as the rapidly falling cost of storage, even relative to growth in processing power and bandwidth. This creates an opportunity for increased transparency through pervasive data collection and monitoring. For example, he notes that companies who invest in security awareness training may be motivated to collect data for tracking violations of data policies. Such increased transparency might enable companies to measure the effectiveness of training, ultimately driving down the number of data policy violations.

Stuart Schechter described an effort to increase the openness of the labor market for identifying vulnerabilities in "Vulnerability Hunters: Surveying Participants in a Poorly Understood Labor Market". Vulnerability markets try to differentiate the security levels of software by establishing a market price for undisclosed vulnerabilities in different types of software. While quasi-markets for vulnerabilities do exist, data regarding their operation are very opaque. Prices paid are not published, and little is understood about vulnerability hunters themselves. To answer these questions, Schechter has devised a survey which he plans to issue to contributers to the Open Source Vulnerability Database.

In "Security through Obscurity: When it Works and When it Doesn't", Peter Swire studies the case for transparency in the disclosure of vulnerabilities and attacks on computer systems. He notes that open source software developers and military types take diametrically opposed views on whether to disclose weaknesses. Swire then argues that each can be correct depending on the circumstances. Disclosure is helpful when it primarily assists defenders without helping attackers much, while disclosure is bad whenever attackers stand to gain more than defenders from the information. Swire suggests that disclosure is often more appropriate for attacks on computer systems because defense mechanisms like firewalls and encryption algorithms do not benefit from hiddenness. Attackers rely on vulnerability information kept secret from defenders because they can subsequently plug the weakness. Swire also discussed the incentives for disclosing security breaches; he finds convergence in the private sector because open source developers often

require some secrecy while proprietary software can benefit from increased openness. However, the government typically has the incentive to inhibit disclosure; thus, freedom-of-information mandatory disclosure is helpful.

Neil Gandal described a model for vulnerability disclosure in "Internet Security, Vulnerability Disclosure, and Software Provision". His model studies the disclosure dilemma facing software vendors: disclosing vulnerabilities and issuing updates protects consumers who install updates, but not all consumers install updates and disclosure can facilitate reverse engineering. The primary question Gandal considers in the model is the effect of mandatory disclosure. Essentially, mandatory disclosure negatively affects "marginal consumers", i.e., those consumers that just value the software enough to purchase it at a lower price. If the firm sets the price of software low enough to attract consumers who do not value patches enough to apply them, then mandatory disclosure can lead to sub-optimal outcomes from a welfare perspective. Gandal also discussed how the model can be used to understand the effect of third parties like CERT or vulnerability markets that can increase vulnerability disclosure rates.

2.3 Economics-informed system design

Economic analysis can prove useful in the design of computer systems and applications, from the legal agreements specifying software use to principles for protecting user privacy.

In "Privacy, Incentives, and Contractual Efficiency in the Market for Consumer Software", Jens Grossklags studies end-user license agreements (EULAs). Regulators are concerned that the terms of EULAs are overly harsh (e.g., they permit adware and spyware). Yet empirical evidence has demonstrated that while consumers differentiate products based on price, they do not consider the terms of the EULA. One common criticism of EULAs is that they are opaque and make it difficult for consumers to bother with studying the terms. Grossklags devises an experiment to discern whether including summaries of terms in EULAs impact the decision on whether to install the software. The study finds that short summaries have a strong impact on user's deciding whether to install software, which suggests that regulatory proposals to mandate such summaries are well-founded.

Anindya Ghose presented techniques for analyzing the economic impact of user-generated product reviews in "Designing Review Ranking Systems: Combining Economics with Opinion Mining". Both buyers and sellers can create web content which influence sales rates. Ghose analyzed Amazon customer reviews and found that more "subjective reviews", i.e., those that

deviated from manufacture specifications, increase sales. Furthermore, even negative reviews can increase sales if they are informative. Ghose argues that devising better rankings of user-generated content and reviews is essential to improving the efficiency of electronic markets.

In "Privacy Engineering", Sarah Spiekermann distinguished between "privacy by policy" and "privacy by architecture". She disparaged privacy by policy as a shortcut approach where companies issue fair information practices pledging to not abuse personal information while continuing to collect personally-identifiable data. Privacy by architecture, by contrast, protects privacy through system design choices. She uses the location-based services as an example. These could be provided at the network level if the network provider is made aware of device location. But these services could also be provided as a client-based solution where the network operator need not be aware of the client device's location.

In "Fuzzy MLS: An Experiment on Quantified Risk-Adaptive Access Control", Pau-Chen Cheng proposed a mechanism for explicitly factoring risk into access control decisions. The work is motivated by the widespread proliferation of ad-hoc exceptions to access control policies. Instead, using Cheng's system access control decisions are made by explicitly weighing potential costs of breaches against the benefit of access. They adopted their scheme to a standard multi-level access control policy by turning access control decisions into risk-weighted calculations based on user input rather than binary decisions.

Several presenters described computer system applications in need of an incentive-based solution. In "Routing Security Economics", Stephen Bellovin described weaknesses in Internet routing protocols that can be abused by a dishonest network participant. He noted that these routing attacks should be distinguished from software vulnerabilities or buggy code; instead, the security of Internet routing depends on the exclusion of malicious behavior by Internet service providers (ISPs). He notes that the costs of replacing these protocols with more secure ones (i.e., protocols that require authentication in routing advertisements) is high enough to hinder their adoption. Instead, some large ISPs have deployed mitigation strategies like deaggregation that are less expensive to the originator but place additional burden on everyone else. Consequently, Bellovin has identified routing security as a problem in need of an economic solution.

In "Countermeasures against Government-Scale Monetary Forgeries", Nicolas Christin outlined a solution using barcodes and online verification to detect even near-perfect cash forgeries. The trouble with existing cash is that it can be usefully duplicated by an adversary with government-level

resources (i.e., enough resources to build a comparable printing press). Furthermore, it can be difficult to verify whether the cash is uniquely legitimate. Adding a 2-d barcode makes cash universally verifiable; to prevent duplication, he described an online lookup protocol that can be initiated between a consumer and the bank. Christin argued that while the addition of a small amount of forgeries to the money supply has a minimal macroeconomic impact, stopping forgeries is nonetheless important since an adversary could cause local destabilization to the money supply and since forged money could be used on the black market for nefarious purposes. He also noted that the forgery arms race is distinguished from other types because the defender can easily stop the attacker by incorporating his proposed solution.

In "Design of a blocking-resistant anonymity system", Roger Dingledine discussed a potential solution to one aspect of the arms race between the designers of Tor, an anonymous communication system, and censoring governments. Tor relies on volunteer servers to route its users' traffic. At present, these servers address information is publicly disclosed. A censoring government can easily block this information if it desires. Dingledine discussed a countermeasure called bridges, where regular Tor users distribute routing information. Dingledine then described several techniques to protect bridges themselves, from releasing certain bridges during different time periods to distributed the bridge details over a social network.

In "Valet Services: Improving Hidden Servers with a Personal Touch", Paul Syverson considers the same problem of hiding identifying information from adversaries while still providing services. Here, the aim is to design a protocol that minimizes identifiable information tied to volunteer nodes to make it harder for attackers to target them. Syverson describes the use of valet nodes which hide service introduction points. These valet nodes can also be used to differentiate the quality of service provided to system participants.

2.4 Rational Security Investments

In "Modelling and Economics of IT Risk Management and Insurance", Costas Lambrinoudakis presents a Markov model for security investment in the face of uncertainty over attacks. The model requires accurate transition probabilities and loss estimates, which Lambrinoudakis admits may be difficult to obtain. The model is used to calculate optimal security investment and optimal insurance contracts. Essentially, resources are devoted to threats based upon the potential harm and probability of occurrence.

In "Models and Measures for Correlation in Cyber-Insurance", Gaurav

Kataria discussed the problem of interdependent risk and its effect on the market for cyber insurance. Firms' IT infrastructure is connected to other entities – so its efforts may be undermined by failures elsewhere. Cyber attacks also often exploit a vulnerability in a program used by many firms. Interdependence can make some cyber-risks unattractive to insurers – particularly those risks that are globally rather than locally correlated, such as worm and virus attacks, and systemic risks such as Y2K. Kataria argued that risks with high internal correlation but low global correlation, such as insider attacks, provide the best opportunity for cyber insurance.

Sometimes security measures are advocated even when they overstate benefits without accounting for costs imposed. In "Competing with Free: The Impact of Movie Broadcasts on DVD Sales and Internet Piracy", Michael Smith examined the justification for anti-piracy measures on high definition broadcast television. He collected data from national movie broadcasts on over-the-air networks and some popular cable networks, sales data through Amazon sales rank and estimated piracy using BitTorrent trackers. He found that movie broadcasts on television have a significant impact on DVD sales (as well as piracy). The impact on DVD sales is strikingly higher: sales increase by 400% initially, compared with a 150% increase in illegal downloading. While Smith was hesitant to make any policy recommendations from the study, he did note that the benefits of broadcasting over television are high enough for movie studios to consider making watching movies over television simple to maximize viewership.

2.5 Network economics

Several presenters portrayed networks as an emerging paradigm for better understanding information security and privacy. In "Surveillance of Emergent Associations: Freedom of Association in a Network Society", Katherine Strandburg discussed the legal implications of network surveillance. She described a world of "emergent associations", where people associate spontaneously using the Internet with wireless and locational technologies. But crucially, these interactions take place via communication intermediaries. Consequently, these spontaneous associations are captured through massive amounts of traffic data (e.g., telephone records, Internet traffic logs and location traces). Indeed, Strandburg noted that highly valuable relational surveillance was possible using social network analysis on non-content traffic data, from investigating a suspicious individual's acquaintances to mining all traffic data for suspicious relational patterns.

Strandburg then described a legal paradox: weak protection of non-

content traffic data yet strong protection of free associations under the US Constitution's first amendment. She then argued that relational surveillance is often used to track associations, so the law should be updated to regulate access to traffic data for network analysis purposes.

In "Network Economics and Security Engineering", Tyler Moore described a framework for modeling repeated attack and defense on networks. He noted that several computing applications, from Internet routing to sensor networks to online social networks, can usefully be represented as a graph of nodes and edges. These networks involve several important characteristics: the distribution of the number of edges each node possesses and the dynamics of adding and removing edges as nodes join, leave and move about a network. He presents a repeated attack and defense scenario where an attacker can remove targeted nodes from a scale-free network, followed by a defender replenishing the network according to different strategies. He showed that naive replenishment fared badly, while replacing removed nodes with localized clique structures is more resilient.

Moore then considered how this repeated attack and defense scenario could be applied to other networks. As an example, he discussed various strategies for punishing misbehavior in a distributed wireless network. One is to allow devices to vote for a misbehaving node's removal. An active attacker might try to disrupt the network by voting against honest nodes. Under a repeated attack and defense framework, attacker nodes vote against honest nodes followed by defenders attempting to remove bad nodes. Finally strategies are updated by nodes deciding whether to alter the threshold required to remove bad nodes if enough unpunished nodes remain.

In "Network formation, Sybil Attacks and Reputation Systems", George Danezis modeled how Sybil attacks (where malicious devices pretend to be many different identities) impact the formation of peer-to-peer networks. He starts with a simple game where nodes wish to connect to their friends. Nodes have a finite link budget which ensures that they cannot directly connect to all of their friends. Node utility is a negative sum of the length of shortest paths to all friends. Danezis remarked that modeling network formation game theoretically proved significantly more complicated than he initially expected. To keep the model simple, he allowed nodes only two choices of strategy : nodes connect only to friends, or nodes use half of their links to connect to friends and the other half to connect to strangers. The only Nash equilibrium found thus far is for everyone to only connect to friends. However, he does not believe this is necessarily the only one. Danezis concluded by expressing frustration in the lack of power and realism of game theory in modeling computer networks.

3 Breakout sessions

Participants were divided into four small groups to debate a research agenda for information security economics. Participants were encouraged to discuss their chosen methodology, as well as pose both open and closed research questions.

As the workshop participants had diverse backgrounds, the methodologies described were correspondingly varied. Some found human subject experimentation to be a promising method for understanding how people actually respond to computer systems as well as incentives. Similarly, participants emphasized the need to complement theoretical models with user studies. This led to the complaint that many existing economic models of information security are too simplified to be useful. Participants also found it difficult to decide under what circumstances to use different modeling techniques (e.g., input-output models vs. game-theoretic ones).

Before enumerating the many open questions in security economics, participants described several key closed questions. First is that security always presents trade-offs, even when these trade-offs are unclear. A corollary to this is that security is a cost center rather than a profit center. Information security does not have an exclusively technical solution.

Many open questions remain. What is the optimal level of security investment? There is a clear disconnect between theoretical models of security investment and actual business practices. What is the proper role of government in enforcing investment? Can the lightweight framework of mandating increased transparency be usefully applied to other areas besides privacy breaches? Under what circumstances to people value privacy? Is the claim that users do not care about privacy actually a fallacy? To what degree are user perceptions of security skewed by vendor marketing? In many circumstances, quantifying the benefits of security and privacy remains elusive. Does this mean there is more research to be done, or are some benefits (and costs) simply unquantifiable?

Many security threats have very low likelihood of occurring based on past occurrences even if the technical potential is very real (e.g., routing attacks, attacks on critical infrastructures like SCADA systems and the telephone network). What is the appropriate security investment for protecting against devastating attacks that might occur with very low probability?

4 Conclusions

Over the past several years, a research program on the economics of information security has built many cross-disciplinary links and has produced many useful insights from unexpected places. Many perverse aspects of information security that had been long known to practitioners turn out to be quite explicable in terms of the incentives and market failure. The DIMACS workshop brought together researchers active in this field to discuss future directions.

The workshop established a need to improve modeling to better understand user behavior. In particular, psychological explanations of risk management may prove helpful; user studies are necessary to understand how technical solutions are used in practice; and human subject experimentation can be used to evaluate trade-offs in system design.

Workshop participants also emphasized the need to design systems to satisfy incentives. Participants identified applications that could benefit from economic analysis, from user-generated product reviews to privacy-preserving technologies to mitigating weaknesses in Internet routing protocols.

Finally, workshop participants identified the need for increased transparency in the provision of information security. Speakers analyzed the factors supporting and inhibiting mandatory privacy-breach notifications, as well as presenting a method for indirectly measuring its impact on security investment for different industries. The merits of transparency through vulnerability disclosure were also debated, from presenting models of its impact on social welfare to identifying circumstances where disclosure is harmful to surveying the vulnerability hunters themselves. While there are many subtleties to be studied in determining appropriate disclosure, the overall conclusion was that more transparency is needed throughout.

5 Acknowledgments

The organizers and author would like to thank the Center for Discrete Mathematics and Theoretical Computer Science for the convocation of the workshop. The organizers would like to thank the staff of DIMACS, the notetakers, and presenters. The organizers also wish to thank the I3P.