ECONOMICS OF INFORMATION SECURITY

Advances in Information Security

Sushil Jajodia
Consulting Editor, Center for Secure Information Systems
George Mason University, Fairfax, VA 22030-4444
email: jaiodia@gmu.edu

The goals of Kluwer International Series on ADVANCES IN INFORMATION SECURITY are one, to establish the state of the art of, and set the course for, future research in information security and two, to serve as a central reference source for advanced and timely topics in information security research and development. The scope of this series includes all aspects of computer and network security and related areas such as fault tolerance and software assurance.

ADVANCES IN INFORMATION SECURITY aims to publish thorough and cohesive overviews of specific topics in information security, as well as works that are larger in scope or that contain more detailed background information than can be accommodated in shorter survey articles. The series also serves as a forum for topics that may not have reached a level of maturity to warrant a comprehensive textbook treatment. Researchers as well as developers are encouraged to contact Professor Sushil Jajodia with ideas for books under this series.

Additional titles in the series:

PRIMALITY TESTING AND INTEGER FACTORIZATION IN PUBLIC KEY CRYPTOGRAPHY by Song Y. Yan; ISBN: 1-4020-7649-5

SYNCHRONIZING E-SECURITY by Godfried B. Williams; ISBN: 1-4020-7646-0

INTRUSION DETECTION IN DISTRIBUTED SYSTEMS:
An Abstraction-Based Approach by Peng Ning, Sushil Jajodia and X. Sean Wang ISBN: 1-4020-7624-X

DISSEMINATING SECURITY UPDATES AT INTERNET SCALE by Jun Li, Peter Reiher, Gerald J. Popek; ISBN: 1-4020-7305-4

SECURE ELECTRONIC VOTING by Dimitris A. Gritzalis; ISBN: 1-4020-7301-1

APPLICATIONS OF DATA MINING IN COMPUTER SECURITY, edited by Daniel Barbará, Sushil Jajodia; ISBN: 1-4020-7054-3

MOBILE COMPUTATION WITH FUNCTIONS by Zeliha Dilsun **Kırlı,** ISBN: 1-4020-7024-1

TRUSTED RECOVERY AND DEFENSIVE INFORMATION WARFARE by Peng Liu and Sushil Jajodia, ISBN: 0-7923-7572-6

RECENT ADVANCES IN RSA CRYPTOGRAPHY by Stefan Katzenbeisser, ISBN: 0-7923-7438-X

E-COMMERCE SECURITY AND PRIVACY by Anup K. Ghosh, ISBN: 0-7923-7399-5

INFORMATION HIDING: Steganography and Watermarking-Attacks and *Countermeasures* by Neil F. Johnson, Zoran Duric, and Sushil Jajodia, ISBN: 0-7923-7204-2

Additional information about this series can be obtained from
http://www.wkap.nl/prod/s/ADIS

ECONOMICS OF INFORMATION SECURITY

edited by

L. Jean Camp
Harvard University, U.S.A.

Stephen Lewis
University of Cambridge, UK

KLUWER ACADEMIC PUBLISHERS
NEW YORK, BOSTON, DORDRECHT, LONDON, MOSCOW

eBook ISBN: 1-4020-8090-5
Print ISBN: 1-4020-8089-1

Created in the United States of America

Visit Springer's eBookstore at: http://www.ebooks.kluweronline.com
and the Springer Global Website Online at: http://www.springeronline.com

Contents

Preface

The security market has failed.

On Tuesday, October 8, 2003 Aaron Caffrey, age nineteen, began his trial. The charge: subverting the operation of the Port of Houston. His prosecution had been a model of international interaction, with the British and American authorities cooperating at every step. Mr. Caffery was to be tried in the United Kingdom.

The Port of Houston took all normal security practices. The Port had developed web-based services for assisting shipping pilots as they moor, in coordinating loading and unloading companies, and in harbor navigation. In a denial of service attack Aaron brought the port to a halt on September 20, 2001. (A denial of service attack consists of repeated initiations of contact, with the attacking machine pretending to be many different machines. An analogous attack would be to repeatedly call someone on the phone and remaining silent until the hearer hangs up, then repeating the process constantly so no work could be completed.) The initial stated reason for the attack? A person from Houston had taunted Aaron about the object of his on-line affections.

Aaron Caffery walked free from that courtroom in October 2003. Security experts explained that there was no way to disprove his assertion that his threats against Houston, his association with a hacker group, and his talents proved nothing. The defense illustrated that there was no way to illustrate beyond a reasonable doubt that Caffery's machine itself was not subverted, so that it acted upon direction other than its owners.

A hacker who can both manipulate code and illustrate that no one is immune to hackers, Aaron Caffrey is an autistic young man.

This is the state of the security of the American information infrastructure.

In July, 2003 a virus, a variant of one originally named SoBig, infected one out of every three computers in China. The virus provides spammers with the processing power and bandwidth of the infected computer in their distribution of unwanted mass email. The virus caused mail server crashes, denial of service attacks, and encouraged the spread of an unrelated virus masquerading as a Microsoft patch for SoBig. SoBig was the most expensive in history – until MyDoom arrived six months later. In

the time it takes to publish this work, another even more virulent and expensive virus will undoubtedly appear.

This is the state of the security of the global information infrastructure.

Certainly, the web server at the Port of Houston was economically and politically important enough to warrant sufficient investment in security. Indeed, the Port of Houston is important enough that a single teenager should not be able to single-handedly stop the port from functioning.

Similarly, the investment in personnel, networks, and sheer mass of individual time would argue that a virus such as SoBig would have been more effectively prevented than battled, or tolerated as a chronic insolvable problem, like malaria in the tropics.

Why have market mechanisms thus far failed to create secure networks?

The Internet is critical to all sectors of the economy and integrated into government. Security technologies do exist, and capable programmers can implement secure code. Programming projects and operating systems based on secure design principles populate research databases. Yet the network at the Port of Houston was sabotaged by a creative teenager with limited programming experience.

Why? Clearly the answer to this question must include more than technology. There is a problem in the economics of security, and more broadly in the economics of information control. These problems emerge as security violations, spam, 'private' databases indexed by Google, and products based on practices exposed as snake oil decades before.

Computer viruses and worms are no longer the domains of experts only. Every business experienced infections and disruptions from infected machines in the latest generation of worms. Economics combined with a management, organization theory, and computer security together can address the chronic problems of economic security. Yet the problems of security have not, before now, been systematically examined in economic and management terms. This text, rather than trying to encourage managers and practitioners to become security experts uses the tools of economics to bear on the problems of network security. The result is a narrative about the economic problems of information security, a set of tools for examining appropriate investment in computer security, all embedded in a set of rich metaphors for balancing the various alternative for computer security.

The security market in the case of networked information systems can be thought of in many different ways, and each view suggest a different set of regulatory and economic responses. Yet, for all the metaphors that may apply there is a single potential measure: dollars. Economics offers a powerful lens for understanding the apparently wildly irrational behavior of software providers, companies, home users and even nation states. This text brings all the tools of economics to bear on the individual, corporate, and national problems of computer security. Perverse

incentives, lock-in, irrational risk evaluations and bad information all play a role in creating the chronically broken network.

The economics of information security is not a metaphor for computer security, like war or health. Recognizing the economics of information security allow managers to alter incentives and policy makers to better evaluate policies that may be presented under the warfare metaphor.

A simple example of corporate incentives is that of patching vulnerabilities. Individual departments must pay for their own IT services, machines, and employee time. Engaging ITS to support employees and requiring employees to patch creates immediate costs for each manager. Charging each section for vulnerabilities will enhance company wide security, but such a solution comes from consideration of the complexities of the security market. Assuming that security works like all other goods has and will continue to result in the creation of perverse incentives that cause managers to ignore the long term issue of security in favor of goals with more pressing time frames.

While the elephant of computer security emerges piecewise, with the ear and tail and foot, the volume as a whole offers a clear picture of computer and information security. Such clarity could only be obtained by painting the whole picture with the palette provide by economics.

Camp's article discusses the concept of security vulnerabilities as an externality, and the direct implication of such externalities for market construction. Of course the use of economics proposes that security must be some kind of tradable or measurable good. Perhaps security is that canonical economic failure – a public good. In this case one person's security investment is another's gain, therefore no one makes the adequate investment. Or perhaps it is not the value to others but the simple lack of return that means that there is little investment. If security is an externality it can still be subject to measurement. Understanding security as an externality may inform the security debate and, as the chapter concludes, offer some insight in how to manage it in a corporate environment.

Yet perhaps vulnerabilities and externalities is too narrow a description of security. What kind of good exactly is being measured? Hal Varian offers three scenarios.

First, security can be defined by the lowest investment, just as the height of a protective wall is defined by its lowest or weakest point. Even barbarians knew this, as they aimed for the gate and not the towers.

Second, the level of security can be determined by the greatest investment, as when the town is protected by concentric walls. The highest wall provides the greatest protection (or rather, the combination of the strongest gate and highest wall).

Alternatively, the security level can be determined by the average investment. In this case consider the community involved in the construction of the wall – the wall is as high as the combined effort of all participants. Individual effort can raise the average somewhat, but not

significantly raise the wall. Consumer behavior reflects the assertion that security and privacy claims are not trustworthy. Few consumers exhibit the understanding of "trusted" computing as trustworthy. Indeed, security is more complex than most goods in that its primary function will be subverted by its users. Passwords written on post-it notes, shared passwords, violations of security policy, and sharing of security information are all common. Why is security both so desirable and so frequently subverted?

Control and verification of information are the critical goal of security and privacy. Yet control of information on an individual machine may be of interest to more then the user. In the most common examples, a remote party with commercial interests will want to constrain the use of information; however, even more common is the desire of en employer to control information use on the employee's machine. One economics of security is needed to analyze remote control of information, whereas distinct economic concepts are required to discuss the protection of a set of machines with a define periphery.

Digital rights management systems are designed by producers with complex commercial interests; these interests are often in conflict with the interests of the user. As a result, the most consistent and highest investment in security has been in the interest of manufacturers, not consumers. Trusted computing has been primarily used to implement bundling. Cell phone companies tie the battery to the phone; automobile companies tie maintenance to the dealership. What would be theoretically prevented in the contract can be prohibited by the code.

Ross Anderson has illustrated this dichotomy in a series of case studies of security as applied in modern technologies. The nature of security as a good is complicated by the fact that it is inherently a bundled good. You cannot purchase security in the abstract. There must be a threat to be considered and the security investment (average, lowest or highest) must be commensurate with and targeted to that threat. In all of these the threat as perceived by the user is the threat of external control; while the threat as perceived by the producer is that of a consumer out of control.

Having acknowledged that producer security is at odds with consumer Desires, it is feasible to examine investment from the perspective of the producer or the consumer. Beginning with the producer, Stephen Lewis asks if producers have accurately and correctly invested in digital rights management technology. Indeed, as shown in the next chapter by Stuart Schechter, investments in encryption against P2P networks are in fact changing the balance. But the balance is being changed in favor of the file traders and against the interests of those who would license the content. Beginning with the argument about the current uses of security technology, observing the incentives in peer to peer systems, the final chapter in this section argues that trusted computing may end up supporting the user and subverting the investors.

Indeed if reliable security information is so difficult to find, the incentives so hard to evaluate, and the results so unreliable, why should anyone share it? What are the economic consequences of sharing information? Esther Gal-or and Anindya Ghose examine the generic question of sharing security information, to find that it is in fact anything but generic. The size of the firm, the nature of the market in which the firm is competing, and even the functional requirements for anti-trust policy. Information sharing among firms and across industries varies widely, and this chapter explains why.

Hussein offers a broad look of the quantitative examinations of computer security economics. The findings are remarkably consistent for a young branch of the dismal science. There are a few discordant findings, illustrating that there is no single unified theory of information security but that a range of possibilities suggests reasons for underinvestment.

If security and confidentiality are primarily targeted at preventing firm loss, then what are the limits to security? If security is primarily a conceptual issue, then attacks on reputation as well as integrity are a security issue. Considering the vast investment in brands, are investments in security rational?

Sharing information may lead to more investment and thus a decrease in losses to security breaches. Beyond direct loss, what is the loss in value of the firm when there are security breaches? Larry Gordon and Marty Loeb illustrate that security breaches by and large have little effect on stock market evaluation of a firm. Yet when confidentiality is lost, then there is a high price to pay. The implicit argument is that the market responds very strongly to losses of privacy and less strongly to losses of security. The security market cannot be extricated from the privacy market, without serious misunderstandings of both.

In rejecting techniques that require effort, users are rejecting investment in the very confidentiality that the market so values. Aquisti argues that is because users share the characteristic so often identified in the stock market itself: extremely high long term discounting. Users value the current convenience offered by privacy violations at current value, and implement extraordinary discounts for the later potential harm.

This observation is validated from an entirely different perspective by Paul Syverson in his examination of the security market. Discounts and probabilities are not well understood when consumers offer information that could be used against them. However that immediate discount is extremely well understood.

Shostack makes a counter observation that it is perhaps not the discounts and risk calculations that make users so casual about protecting their own information. Perhaps users simply have no understanding of the threat. Just as some miners refused to take the accumulation of gas seriously as a threat, and no one understood why workers on the Brooklyn Bridge were dying of the bends, individuals today do not understand the value of privacy. To make an analogy, why would someone buy cur-

tains and then offer details of their home over the Internet? The value of security for the end user is even more difficult to understand than the value of privacy for the consumer. The overall evaluation of the security market when seen from the privacy perspective is not optimistic.

Landwehr argues explicitly that the information flows in the security market are broken. Not only do consumers not understand the issues of privacy and security risks, but even vendors themselves do not understand security. Bill Gates' vaulted commitment to security includes training in security for 7,000 developers, yet there has not been a month without the release of a security patch for Microsoft. Even the considerable financial and technical resources of Microsoft cannot result in coherent application of security research implemented decades ago in a complex computing environment characterized by unpredictable interactions.

If security and privacy policies are "lemons markets", then simple claims of investment in security are far cheaper and easier than actually securing a site. If the claims are security are adequate to insure customer trust (and possibly cause malevolent profit-oriented actors to target others) then there is no reason for investment in security or privacy. Like false claims about a reliable used cars, false claims of secure software and false claims of privacy policy have no costs. Ironically, the lemons argument suggest that the core security failure in the information infrastructure is one of trustworthy information. Vila and Greenstadt argue clearly for this counter-intuitive possibility.

Integrating personal actions in security and privacy is a significant contribution of the next chapter. SoBig, MyDoom, and many other viral variants depend on a large population of unsecured user machines to flourish. Users express great concern for security, and privacy concerns have been monotonically increasing. Given this concern, how can observed user behaviors that illustrate that users share information readily and avoid installing security patches be explained?

Acquisiti uses the issue of on-line and off-line identities to illustrate how economics can shed light on the apparent irrationalities of both individuals and the market, regarding the confidentiality of information.

Odlyzko explains that users are correct in rejecting security designed for them by merchants and providers because the greatest value for merchants in controlling information is to implement price discrimination. Offering information to a merchant who can then charge you more is not in the interest of a consumer, even if the issues of control were not relevant. Security systems that violate privacy are directly opposed to the interest of the user when price discrimination is more likely than personal security loss. In economic terms, users are balancing risks when selecting privacy.

A more detailed discussion of users who reject security is provided in the aptly-titled, "We Want Security But We Hate It: The Foundations of Security Techo-Economics in the Social World". The undercurrents

of user resistance to security include economics, as well as being a social and psychological phenomena. Beyond losing money through price discrimination, users seek to maintain control and confidentiality. When much security is implemented in order to best reflect vendor needs (as when security is provided as part of digital rights management) users seek to avoid the "features" offered in mainstream security solutions.

Perhaps users are motivated but misinformed. Certainly, corporate organizations are not discouraged from investing in security because of concerns of control of the desktop - this would be a feature and not a bug. Perhaps the critical problem in the information age is the information flow. Information is calculated and generated. Standards are made. Committees meet. Yet for all the research and effort, homes users do not see themselves at risk. Corporations do not develop appropriate responses.

In fact, manipulation of information and users remains a threat that cannot be addressed through technology alone. Can economics hope to address the problems of manipulation of authorized individuals and naive home users? Economics and markets themselves can be manipulated with the same tools of misinformation. "Cognitive Hacking" can apply to economic systems and information systems.

Yet within the generally bleak picture of information failure, market failure and suspicion there are cases of remarkable success. We end with two of these: secure sockets layer and the cable industry.

Having used economics to extract the distinctions between security and privacy as information control mechanisms in the market, the book closes with some specific examples of security in markets.

The story of the secure sockets layer and secure telnet illustrate that a chronic low level of security need not be an external state of affairs, no matter how long term or ubiquitous the state of affairs. The cable industry illustrates that lock-in need not lock out security, if the incentives are properly aligned. The following examination of the secure shell and the secure sockets layer illustrates that forward movement is possible even in a distributed, chaotic market. However, even the success stories of Larochelle and Rosasco illustrate that history offers as much caution as promise, as each tale offers specific conditions and constraints that enable security diffusion.

Economics offers a powerful lens for the examination of security. This text aims to promote a more sophisticated vision of security in an effort to assist designers in making systems that respect the alignment of incentives, managers in aligning their investments with the most critical security problems, and policy makers in understanding the nature of the chronic, core problem of modern computer security. Bruce Schneier explains better than any how apparently technical failures are in fact economic failures, and his explanation provides the final thoughts in this text.

Incentives in the security market are badly aligned, and the technology is not understood. Ironically in the information age, trustworthy information is increasingly difficult to locate. To paraphrase Mark Twain: A virus can be half way around the world while a patch is still putting its boots on.

L JEAN CAMP

Acknowledgments

Certainly the most obviously deserving of acknowledgement in this volume are the contributors. And indeed they are deserving. We are also in the debt of our editors, Sharon Palleschi and Susan Lagerstrom-Fife. Thank you for your guidance and support.

Individually, Jean Camp has many to acknowledge.

I would like to first and foremost thank my doctoral candidates who have read and discussed this content with much patience. Warigia Bowman and Allan Friedman, I am in your debt. Having now some experience as an advisor makes me far more grateful for those on my own committee. My advisors, Marvin Sirbu and Doug Tygar, were important sources of critical support in my own intellectual path, one that has crossed many disciplinary lines. As for my other committee members, Pamela Samuelson, Mary Shaw and Granger Morgan, the farther I am from them the larger I see that they are. As with so many things, I did not appreciate them then as much as I now know to admire them.

Ross Anderson and Hal Varian have lead the creation of a field by the organizing of interdisciplinary workshops. Before the First and Second workshops on the Economics of Information Security, most of us were breaking down individual barriers in our own corners of the academy: business, mathematics, electrical engineering, economics, business and public policy. With the organization of a series of workshops, and now the publication of this text based loosely on those workshops, an new arena of discourse has formed. We might not have stormed the barricades, but they are certainly sufficiently honey-combed to be highly permeable.

Andrew Odlyzko also deserves unique mention in my own life as a scholar. Ceaselessly a genial academic gentleman, he has offered guidance to more than one junior faculty member and I am lucky to be among them.

Stephen Lewis would like to thank all of his friends and colleagues both within the Computer Laboratory and outside; a special debt of gratitude is owed to Ross Anderson, for his endless support, guidance and patience as a supervisor.