THE LAW AND ECONOMICS OF Cybersecurity

Edited by
Mark F. Grady and
Francesco Parisi

CAMBRIDGE

CAMBRIDGE www.cambridge.org/9780521855273

This page intentionally left blank

THE LAW AND ECONOMICS OF CYBERSECURITY

Cybersecurity is a leading national problem for which the market may fail to produce a solution. The ultimate source of the problem is that computer owners lack adequate incentives to invest in security because they bear fully the costs of their security precautions but share the benefits with their network partners. In a world of positive transaction costs, individuals often select less than optimal security levels. The problem is compounded because the insecure networks extend far beyond the regulatory jurisdiction of any one nation or even coalition of nations. This book brings together the views of leading law and economics scholars on the nature of the cybersecurity problem and possible solutions to it. Many of these solutions are market based, but they need some help, either from government or industry groups, or both. Indeed, the cybersecurity problem prefigures a host of 21st-century problems created by information technology and the globalization of markets.

Mark F. Grady is Professor of Law and Director of the Center for Law and Economics at the University of California at Los Angeles School of Law. He specializes in law and economics, torts, antitrust, and intellectual property. He received his A.B. degree summa cum laude in economics and his J.D. from UCLA. Before beginning his academic career, Grady worked for the Federal Trade Commission, the U.S. Senate Judiciary Committee, and American Management Systems.

Francesco Parisi is Professor of Law and Director of the Law and Economics Program at George Mason University School of Law and Distinguished Professor of Law at the University of Milan.

THE LAW AND ECONOMICS OF CYBERSECURITY

Edited by

Mark F. Grady
UCLA School of Law

Francesco Parisi
George Mason University School of Law

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 2RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521855273

First published in print format 2005

ISBN-13 978-0-511-13830-0 eBook (Adobe Reader)
ISBN-10 0-511-13830-X eBook (Adobe Reader)

ISBN-13 978-0-521-85527-3 hardback
ISBN-10 0-521-85527-6 hardback

CONTENTS

Acknowledgments *page* vii
Contributors viii

The Law and Economics of Cybersecurity: An Introduction 1
Mark Grady and Francesco Parisi

PART ONE: PROBLEMS

Cybersecurity and Its Problems

1 Private versus Social Incentives in Cybersecurity: Law and Economics 13
Bruce K. Kobayashi

2 A Model for When Disclosure Helps Security: What Is Different about Computer and Network Security? 29
Peter P. Swire

Intervention Strategies: Redundancy, Diversity and Autarchy

3 Peer Production of Survivable Critical Infrastructures 73
Yochai Benkler

4 Cybersecurity: Of Heterogeneity and Autarky 115
Randal C. Picker

PART TWO: SOLUTIONS

Private Ordering Solutions

5 Network Responses to Network Threats: The Evolution into Private Cybersecurity Associations 143
Amitai Aviram

6 The Dark Side of Private Ordering: The Network/Community Harm of Crime 193
Neal K. Katyal

Regulation and Jurisdiction for Global Cybersecurity

7 Holding Internet Service Providers Accountable 221
Doug Lichtman and Eric P. Posner

8 Global Cyberterrorism, Jurisdiction, and International Organization 259
Joel P. Trachtman

Index 297

ACKNOWLEDGMENTS

The editors of this volume owe a debt of gratitude to many friends and colleagues who have contributed to this project at different stages of its development. Most notably, we would like to thank Emily Frey, Amitai Aviram, and Fred Wintrich for encouraging and helping coordinate the planning of this project. The Critical Infrastructure Protection Project and the George Mason University Tech Center provided generous funding for the Conference on the Law and Economics of Cyber Security, which was held at George Mason University on June 11, 2004. At this conference, several of the papers contained in this volume were originally presented. David Lord scrupulously assisted the editors in the preparation of the manuscript for publication and in the drafting of the introduction. Without his help, this project would not have been possible. Finally we would like to thank University of Chicago Press for granting the permission to publish the paper by Doug Lichtman and Eric Posner, which will appear in Volume 14 of the *Supreme Court Economic Review* (2006).

CONTRIBUTORS

Amitai Aviram Assistant Professor of Law, Florida Sate University College of Law

Yochai Benkler Professor of Law, Yale Law School

Mark Grady Professor of Law, University of California at Los Angeles, School of Law

Neal K. Katyal John Carroll Research Professor, Georgetown University Law Center

Bruce K. Kobayashi Professor of Law and Associate Dean for Academic Affairs, George Mason University School of Law

Doug Lichtman Professor of Law, University of Chicago Law School

Francesco Parisi Professor of Law and Director, Law and Economics Program, George Mason University School of Law

Randal C. Picker Paul and Theo Leffmann Professor of Commercial Law, University of Chicago Law School; Senior Fellow, The Computational Institute of the University of Chicago and Argonne National Laboratory

Eric P. Posner Kirkland and Ellis Professor of Law, University of Chicago Law School

Peter P. Swire Professor of Law and John Glenn Research Scholar in Public Policy Research, Ohio State University, Moritz College of Law

Joel P. Trachtman Professor of International Law, Fletcher School of Law and Diplomacy, Tufts University

THE LAW AND ECONOMICS OF CYBERSECURITY: AN INTRODUCTION

Mark Grady and Francesco Parisi

Cybercrime imposes a large cost on our economy and is highly resistant to the usual methods of prevention and deterrence. Businesses spent about $8.75 billion to exterminate the infamous Love Bug. Perhaps far more important are the hidden costs of self-protection and losses from service interruption.

Unlike traditional crime, which terrorizes all but has far fewer direct victims, cybercrime impacts the lives of virtually all citizens and almost every company. The Computer Security Institute and the FBI recently released the results of a study of 538 companies, government agencies, and financial institutions. Eighty-five percent of the respondents reported having security breaches, and 64% experienced financial loss as a result (Hatcher 2001). Because this problem is growing on a daily basis, it is imperative that society identify the most economically efficient way of fighting cybercrime. In this volume, the authors present a cross section of views that attempt to identify the true problems of cybersecurity and present solutions that will help resolve these challenges. In the first section, two authors outline some of the major problems of cybersecurity and explain how the provision of cybersecurity differs from traditional security models.

Bruce Kobayashi examines the optimal level of cybersecurity as compared with traditional security. For example, while it might be more efficient to deter robbery in general, individuals may find it easier to simply put a lock on their door, thus diverting the criminal to a neighbor's house. Although in the general criminal context, the government can act to discourage *ex ante* by implementing a sufficient level of punishment to deter the crime from occurring in the first place, this is not so easily achieved in the world of cybercrime. Because the likelihood of detecting cybercrime is so low, the penalty inflicted would have to be of enormous magnitude to deter it.

In this context, companies can either produce private security goods that will protect their sites by diverting the hacker to someone else or they can produce

a public security good that will deter cybercrime in general. The former route will lead to an overproduction of private security, which is economically inefficient because each company takes individual measures that only protect itself as opposed to acting collectively to stop the cyberattacks in the first place. If collective action is used to produce public security, however, an underproduction will occur because companies will have an incentive to free-ride on the general security produced by others.

Kobayashi suggests using a concept of property rights whereby the security collective can exclude free-riders to eliminate this problem. Since security expenditures are not sufficiently novel or nonobvious to merit protection under patent or copyright law, Kobayashi suggests collective security action supported by contractual restrictions on members.

Peter Swire follows on Kobayahi's basic idea of collective action by introducing the notion of cooperation through disclosure. Swire attempts to answer the question of when disclosure may actually improve security. In probing this question, Swire develops a model for examining the choice between the open source paradigm, which favors disclosure, and the military paradigm, which advocates secrecy. The open source paradigm is based on three presumptions: attackers will learn little or nothing from disclosure, disclosure will prompt designers to improve the design of defenses, and disclosure will prompt other defenders to take action. The military paradigm is based on contrary presumptions: attackers will learn much from the disclosure of vulnerabilities, disclosure will not teach the designers anything significant about improving defenses, and disclosure will not prompt improvements in defense by others. Starting with these two paradigms, Swire offers two further concepts that take a middle ground. The first, the Information Sharing Paradigm, reasons that although attackers will learn a lot from disclosure, the disclosure will prompt more defensive actions by others and will teach designers how to design better systems. For example, the FBI's disclosure of a terrorist "watch list" may enable people to be more attuned to who is a terrorist, but it does so at the cost of alerting terrorists to the fact that they are being scrutinized. Opposed to the information sharing paradigm is the theory of public domain, which holds that although attackers will learn little to nothing from disclosure, disclosure will also not teach designers much and will not prompt many additional security steps by others.

Swire reasons that different scenarios warrant adherence to different security paradigms. Factors such as the number of attacks, the extent to which an attacker learns from previous attacks, and the extent of communication between attackers about their knowledge will influence which model should be followed. In general, secrecy is always more likely to be effective against the

first attack. While this might favor the military paradigm in the realm of physical security because of a low number of attacks and relative lack of communication between attackers, the same assumptions do not necessarily hold true in the realm of cybersecurity. Because cyberattacks can be launched repetitively and at minor expense, secrets will soon be learned and companies will expend inordinate amounts of money vainly attempting to retain their secrecy. Further, as is true in traditional physical security, disclosure can often improve security by diverting an attack, presuming that the level of security is perceived as high.

Swire also argues that there are two specific areas in which the presumptions of the open source paradigm do not hold true. First, private keys, combinations, and passwords should never be disclosed because disclosing them does little to promote security or enhance security design, yet it obviously provides valuable information to attackers. Additionally, Swire argues that surveillance techniques should not be disclosed because an attacker is unlikely to discover them during an attack, and thus in the short run not disclosing them will provide the defender with an additional source of security.

In the second section of Part I, Yochai Benkler argues that cybersecurity is best addressed by making system survivability the primary objective of security measures rather than attempting to create impregnable cyberfortresses. By mobilizing excess capacity that users have on their personal devices, a network-wide, self-healing device could be created. The already existing system of music sharing offers a model for achieving this type of security.

While the sharing of music files is admittedly controversial, the systems that have been put in place to make music sharing a reality offer lessons for how broader cybersecurity can be achieved. Professor Benkler's proposal is based on three characteristics: redundant capacity, geographic and topological diversity, and the capacity for self-organization and self-healing based on a fully distributed system that in nowise depends on a single point that can become the focus of failure. The music-sharing industry has been hit by attacks a number of times, and Napster even had its main center of data search and location shut down. Nonetheless, the data survived because of the above characteristics. File-sharing systems have allowed data and capacity to be transferred to where they are most needed, permitting these systems to survive even after repeated attacks. In many file-sharing systems, because the physical components are owned by end users, there is no network to shut down when it is attacked by cyberterrorism.

This same degree of survivability can also be seen in distributed computing, where it easier for a task to be shared by several computers than to build a single, very fast computer. Benkler concludes his article by looking at different

economic models that suggest when and how the lessons of file sharing can be implemented practically in order to achieve long-term survivability.

The article by Randy Picker examines whether and how security can best be achieved in an industry dominated by one company. Many people have come to believe that market dominance by Microsoft compromises cybersecurity by creating a monoculture, a scenario in which common computer codes help spread viruses easily, software facilities are too integrated and thus lead to security lapses, and software is shipped too soon and thus is not adequately developed to address security needs. In this article, Picker attempts to address these criticisms, believing that they are misdirected and will lead to inefficient results.

Those who believe that the monoculture of Microsoft threatens security often liken the situation to the boll weevil epidemic in the early 1900s. Because farmers in the South cultivated only cotton, when an insect arrived that attacked this crop, their fields and means of livelihood were both devastated. Opponents of monoculture believe that diversification helps insure against loss, whether in agriculture or the world of cybersecurity. Picker points out, however, that one of the primary problems with this logic is that it attempts to deal with the problem from the perspective of supply rather than crafting demand-based solutions. Sure, a farmer can protect against total devastation by diversifying and adding corn as a crop, for example, but if there is no demand for corn, the diversification is futile because consumers will not avail themselves of the corn.

Picker's second criticism of the monoculture theorists is that they argue heterogeneity is the best way to address the massive collapse that can result when a virus invades an interconnected world. However, ensuring that different sectors use different operating systems and computers will not mean that all are protected. When an attack hits, it will still shut down one sector. The only way to provide universal protection would be to have all work done on multiple systems, an inefficient solution to the problem. Picker advocates a security model that is very different from the increased interconnection supported by Benkler. Picker instead advocates autarky, or purposefully severing some of the connections that cause the massive shutdown in the first place. Picker argues that we need to accept the fact that interconnection is not always good. Which is economically more efficient, to have ten connected computers run ten different operating systems or to have ten isolated computers each running Windows?

Picker concludes his article by suggesting that security concerns can be remedied through the use of liability rules. Imposing liability through tort law would, however, create headaches because it would be hard to sort out questions of fault and intervening cause among the developer, the cyberterrorist who unleashed

the virus, and the end user who clicked when he should not have done so. Likewise, requiring the purchase of mandatory insurance would be economically counterproductive. Rather, in Picker's view, partial insurance that focuses on the first wave of consumers who face greater risks (from the less developed product) is the economically most viable solution.

Part II of this volume offers regulatory solutions that address the major problems of cybersecurity. The authors highlight the debate between public and private security by presenting highly divergent positions. Amitai Aviram discusses private ordering achieved through private legal systems (PLSs), institutions that aim to enforce norms when the law fails (i.e., neglects or chooses not to regulate behavior). Aviram's article gives a broad perspective on how PLSs are formed and then suggests practical applications for the field of cybersecurity. Aviram reasons that PLSs cannot spontaneously form because new PLSs often cannot enforce cooperation. This gap occurs because the effectiveness of the enforcement mechanism depends on the provision of benefits by the PLS to its members, a factor that is nonexistent in new PLSs. Thus, new PLSs tend to use existing institutions and regulate norms that are not costly to enforce, ensuring gradual evolution rather than spontaneous formation. PLSs have widely existed throughout history. Literature about PLSs, however, has largely focused on how these organizations develop norms rather than how these organizations come into existence in the first place.

In examining this question, Aviram starts with a basic paradox of PLS formation: in order to secure benefits to its members, a PLS must be able to achieve cooperation, but to achieve cooperation, a PLS must be able to give benefits to its members. This creates a chicken-and-egg situation. While this problem could be resolved through bonding members in a new PLS, bonding is often too expensive. Accordingly, PLSs tend to simply develop and evolve from existing institutions rather than develope spontaneously and independently.

To determine when, how, and by whom a norm can be regulated, it is necessary to understand the cost of enforcing the norm. To understand this, it is necessary to fully comprehend the utility of the norm to the network's members, understand the market structure of the members, and understand what game type and payoffs have been set up by the norm for the network's members. Aviram introduces a variety of gametypes based on the expected payoffs to members. Some of the gametypes have higher enforcement costs, others have lower costs. It is the gametypes that have low enforcement costs that become the building blocks of PLSs, while those with high enforcement costs evolve gradually.

Aviram applies this concept to cybersecurity by looking at networks that aim to facilitate communication and information sharing among private firms.

Unfortunately, these networks have been plagued by the traditional problems of the prisoner's dilemma: members fear cooperation and the divulging of information because of worries about increased liability due to disclosure, the risk of antitrust violations, and the loss of proprietary information. Aviram thinks that part of the reason for the failure of these networks is that they are attempting to regulate norms with high enforcement costs without the background needed to achieve this. Aviram suggests restricting the membership of these networks so that they are not as broadly based as they presently are. This would allow norms to be developed among actors with preexisting business connections that would facilitate enforcement (as opposed to the broad networks that currently exist and cannot enforce disclosure).

The article by Neal Katyal takes a completely divergent position, reasoning that private ordering is insufficient and in many ways undesirable. Katyal argues that we must begin to think of crime not as merely harming an individual and harming the community. If crime is viewed in this light, solutions that favor private ordering seem less beneficial, and public enforcement appears to have more advantages. Katyal maintains that the primary harm to the community from cyberattacks does not necessarily result from the impact on individuals. Indeed, hackers often act only out of curiosity, and some of their attacks do not directly affect the businesses' assets or profits. Rather, these attacks undermine the formation and development of networks. Katyal contends that society can therefore punish computer crimes "even when there is no harm to an individual victim because of the harm in trust to the network. Vigorous enforcement of computer crime prohibitions can help ensure that the network's potential is realized."

Public enforcement is also defended because without governmental action to deter cybercrime only wealthy companies will be able to afford to take the necessary measures to protect themselves. Katyal compares the use of private ordering as the solution for cybercrime to the government's telling individuals that it will no longer prosecute car theft. Indeed, if the government adopted this policy, car theft might decrease because fewer people would drive and those that did drive would take the precautions necessary to protect themselves from theft. While this might seem logical (and has even been used to a large extent in the cyberworld), it fails to take into account exogenous costs. For example, less driving may equal less utility, while the use of private security measures raises distributional concerns (e.g., can only the wealthy afford the security measures necessary to drive?).

Finally, Katyal suggests that to some extent private security measures may increase crime. Imagine a community in which the residents put gates around their homes and bars over their windows. Such measures may deter crime for each individual, but "it suggests that norms of reciprocity have broken down

and that one cannot trust one's neighbor." One result might be that law-abiding citizens would leave the neighborhood, resulting in a higher crime rate. One of the primary reasons for public law enforcement is to put measures into place that are needed to protect the citizens while averting sloppy and ineffective private measures.

Katyal concludes by arguing that not all cybercrimes can be punished and not all should be punished the same way. If the police were to go after every person who committed a cybercrime, it would lead to public panic and further erode the community of trust. Additionally, some crimes, like unleashing a worm in a network, are more serious than a minor cybertrespass.

The article by Lichtman and Posner attempts to move beyond the debate of public versus private enforcement by creating a solution that relies on private measures enforced and promoted by publicly imposed liability. The authors acknowledge that vast security measures have been taken both publicly and privately to address the problem of cybersecurity. However, these measures have not sufficiently addressed the harm caused by cybercrime because the perpetrators are often hard to identify, and even when they are identified, they often lack the resources to compensate their victims. Accordingly, the authors advocate adopting a system that imposes liability on Internet service providers (ISPs) for harm caused by their subscribers. The authors argue that this liability regime is similar to much of tort law, which holds third parties accountable when they can control the actions of judgment-proof tortfeasors. While this idea may run parallel to the common law, the authors acknowledge that it appears to run counter to modern legislation, which aims to shield ISPs from liability. However, even in these laws, the roots of vicarious liability can be seen in the fact that immunity is often tied to an ISP's taking voluntary steps to control the actions of its subscribers.

One of the objections that the authors see to their proposal is related to the problem of private enforcement that Katyal discusses in the previous article. Shielding ISPs from liability, like failing to publicly enforce cybersecurity, will give end users an incentive to develop and implement their own security devices. Lichtman and Posner counter that this argument does not suggest that ISPs should not face liability but that their liability should be tailored to encourage them "to adopt the precautions that they can provide most efficiently, while leaving any remaining precautions to other market actors." Indeed, just as auto drivers are not given immunity from suit based on the argument that pedestrians could avoid accidents by staying at home, the same should hold true in the cyberworld.

The second criticism to this proposal is that it might cause ISPs to overreact by unnecessarily excluding too many innocent but risky subscribers in the name of security. Increased security may indeed drive up costs and drive away marginal

users, but likewise users may be driven away by insecurity in the cyberarena. Posner and Lichtman also believe that the danger of increased cost to ISPs can be alleviated by offering tax breaks to ISPs based on their subscriber base, prohibiting state taxation of Internet transactions, or subsidizing the delivery of Internet access to underserved populations. The problem of viruses traveling across several ISPs can be resolved through joint and several liability, while the fear that no one individual will be harmed enough by cybercrime to bring suit can be resolved through class action lawsuits or suits initiated by a state's attorney general.

The main concern regarding the use of ISP liability is that it would be ineffective because of the global reach of the Internet, for a cybercriminal could simply reroute his or her attack through a country with less stringent security laws. Posner and Lichtman address this concern by arguing that global regimes can be adopted to exclude Internet packets from countries with weak laws. As countries like the United States adopted ISP liability, it would spread to other nations.

Trachtman picks up on this final concern, which is common to many Internet security problems and proposals: the global reach of the Internet and accompanying issues of jurisdiction and international organization. This concern has become even more acute with the development of organized cyberterrorism, as evidenced by the cyberterrorism training camps run by Al Qaeda when the Taliban controlled Afghanistan. Throughout his article, Trachtman examines the same question seen in the articles by Aviram, Katyal, and Posner and Lichtman: to what extent is government regulation necessary to achieve cybersecurity? Trachtman acknowledges that private action suffers to some extent from the inability to exclude free-riders and other collective action problems. Trachtman suggests that private action may be sufficient to resolve some forms of cybercrime, but it clearly will not work to eliminate all cyberterrorism. There are areas that warrant international cooperation, including (1) the limitation of terrorist access to networks, (2) *ex ante* surveillance of networks in order to interdict or repair injury, (3) *ex post* identification and punishment of attackers, and (4) the establishment of more robust networks that can survive attack.

Once it has been decided whether private or public action should be favored, there remains the issue of whether local action is sufficient. Cybercrime proposes unique jurisdictional questions because actions in one country may have effects in another. If the host country will not enforce laws against the cybercriminals, how can the victim country stop the attack? Ambiguous jurisdiction is one of the main problems faced by modern international law in this area. The solution would seem to require international cooperation. Trachtman

suggests creating an umbrella organization that has jurisdiction over these matters and can act transnationally. Trachtman concludes by offering a variety of game theory presentations that exhibit when and how international cooperation can best occur in the realm of cybersecurity.

The authors of the articles in this volume have attempted to provide a resource for better understanding the dilemmas and debates regarding the provision of cybersecurity. Whether cybersecurity is provided through private legal systems or public enforcement or a combination of the two, the development and implementation of new and more efficient tools for fighting cybercrime is high on the list of social priorities.

REFERENCE

Hatcher, Thurston. 2001. Survey: Costs of Computer Security Breaches Soar. *CNN.com.* http://www.cnn.com/2001/TECH/internet/03/12/csi.fbi.hacking.report/.

PART ONE

PROBLEMS

Cybersecurity and Its Problems

ONE

PRIVATE VERSUS SOCIAL INCENTIVES IN CYBERSECURITY: LAW AND ECONOMICS

Bruce H. Kobayashi*

I. INTRODUCTION

Individuals and firms make significant investments in private security. These expenditures cover everything from simple door locks on private homes to elaborate security systems and private security guards. They are in addition to and often complement public law enforcement expenditures. They also differ from public law enforcement expenditures in that they are aimed at the direct prevention or reduction of loss and not necessarily at deterring crime through *ex post* sanctions.[1]

A growing and important subset of private security expenditures are those related to cybersecurity (see Introduction and Chapter 5). Private security expenditures are important given the decentralized nature of the Internet and the difficulties in applying traditional law enforcement techniques to crime and other wealth-transferring activities that take place in cyberspace. These include difficulties in identifying those responsible for cybercrimes, difficulties arising from the large volume and inchoate nature of many of the crimes,[2] and difficulties associated with punishing judgment-proof individuals who are eventually identified as responsible for cyberattacks. As a consequence, those responsible

[1] This analysis does not consider the use of public sanctions and enforcement resources. The level of public enforcement will generally affect the level of private expenditures. For example, public enforcement and sanctions may serve to "crowd out" private expenditures. For analyses of private law enforcement systems, see Becker and Stigler (1974); Landes and Posner (1975); Friedman (1979 and 1984). See also Chapter 7, which discusses the use of vicarious liability as a way to increase security and law enforcement.

[2] For an analysis of punishment for attempts, see Shavell (1990) and Friedman (1991).

*Associate Dean for Academic Affairs and Professor of Law, George Mason University, School of Law. This paper presents, in nonmathematical form, the results presented in Kobayashi (forthcoming). The author would like to thank the Critical Infrastructure Protection Project at George Mason University Law School for funding.

for cyberattacks may perceive both that the probability of punishment is low and that the size of the sanction (when punishment occurs) is small; the final result will be a low expectation of penalty and inadequate deterrence (Becker 1968).

Although individuals and businesses have made significant private investments in cybersecurity, there is a concern that leaving the problem of cybersecurity to the private sector may result in an inadequate level of protection for individuals, firms, and critical networks.[3] Further, private efforts to identify and pursue those responsible for cyberattacks often will redound to the benefit of others, leading to free-riding and inadequate incentives to invest in cybersecurity.[4] This concern has led to calls for government intervention to remedy the perceived underinvestment in cybersecurity.[5]

The purpose of this paper is to examine the basic economics of private cybersecurity expenditures, to examine the potential sources of underinvestment, and to evaluate potential market interventions by the government. This paper begins by reviewing the existing literature on private security expenditures. This literature has concentrated on the provision of private goods such as locks and safes. Such goods are characterized as private goods because, for example, a physical lock or safe protecting a particular asset cannot generally be used by others to protect their assets. In contrast to the perceived underinvestment in cybersecurity, the existing literature does not predict an underinvestment in private security goods. Indeed, the models described in the literature show that, among other things, private security goods may serve to divert crime from protected to unprotected assets and that as a result equilibrium expenditures may exceed socially optimal levels. Further, attempts by firms to reduce wealth

[3] For a discussion of these issues, see Frye (2002). Katyal (Chapter 6) notes the existence of network and community harms caused by crimes that are not internalized by the direct victim of the crime. But see Chapter 5, which notes the benefits of network effects as a mechanism to enforce private norms.

[4] In some cases, firms able to internalize network benefits associated with their products may also be able to internalize the benefits of security expenditures. For example, Microsoft Corporation, in a November 5, 2003, press release, announced the initial $5 million funding of the Anti-Virus Reward Program, which pays bounties for information that leads to the arrest and conviction of those responsible for launching malicious viruses and worms on the Internet. For a discussion of bounties generally, see Becker and Stigler (1974). Microsoft, owing to its large market share, can internalize more of the benefits of private enforcement expenditures. However, its large market share and its de facto status as a standard setter serve to lower the costs of conducting a widespread cyberattack and have resulted in a many attacks directed at computers using Microsoft products. For an analysis of the trade-offs involved with de facto standards in the cybersecurity context, see Chapter 4, which describes the use of decentralized, distributed, and redundant infrastructures as a way to increase system survivability.

[5] Krim (2003) reports that Bush administration officials warn that regulation looms if private companies do not increase private efforts at providing cybersecurity. See also Chapter 6.

transfers that do not represent social costs may also cause private security expenditures to exceed socially optimal levels.

The paper next explores differences between the expenditures on private security goods and expenditures on cybersecurity. It focuses on two primary differences between cybersecurity and the type of security discussed in the existing literature: the public good nature of cybersecurity expenditures and the fact that the social harm caused by a cybercrime greatly exceeds any transfer to the criminal. The paper shows how each of these differences affects the incentives of individuals to invest in cybersecurity. Indeed, both differences serve to reduce any overincentive to invest in private security goods relative to the standard private goods case and suggest an underlying reason why cybersecurity expenditures may be too low. The paper concludes by examining several proposals for government intervention the private market for cybersecurity and how such proposals will address these underlying factors.

II. PRIVATE SECURITY EXPENDITURES

The existing literature on the private provision of security expenditures has focused on cases in which individuals or firms spend resources on private security goods (Shavell 1991).[6] According to the basic model, private individuals invest in goods private security such as locks or safes in order to prevent socially costless transfers. Security goods such as locks and safes are private goods because they cannot be used in a nonrivalrous manner. That is, a lock or safe protecting a particular asset cannot generally be used by others to protect their assets.

In the basic model, criminals expend resources in an attempt to transfer wealth by attacking the sites of potential victims. These potential victims invest in private security to reduce the impact of crime and other wealth-transferring activity. An increase in the level of private security expenditures, *ceteris paribus*, has several primary effects. Additional security expenditures decrease the magnitude of the expected transfer given an intrusion. As a result of this reduction in the expected net gain to the criminal, the equilibrium rate of intrusions will decrease, and the probability of an attack on the assets protected by the security goods will fall.

Under the assumption that the activity addressed by private security expenditures consists of costless wealth transfers, the social objective is to minimize the total resources used by criminals attempting to achieve these transfers and

[6] For an explicit mathematical treatment of this issue, see Kobayashi (forthcoming).

by potential victims attempting to prevent them. The existing literature has identified several reasons that private and social incentives to invest in private security diverge. The first is the existence of uninternalized spillovers or externalities (Shavell 1991; Clotfelter 1978). There are both positive and negative spillovers from private expenditures on security. Positive spillovers include the provision of a general deterrence effect: the expenditures by one individual decrease the net expected gain from wealth-transferring activity, which in turn reduces the general level of criminal activity and thus creates a positive spillover effect for all potential victims. That is, individual expenditures that increase the perceived level of expenditures will protect all sites, including those belonging to individuals who choose to provide no security. This effect will be strongest when criminals know the overall level of private security expenditures but cannot observe individual expenditures until after they have made an effort to engage in a wealth transfer. More generally, even observable private security goods, through their tendency to reduce the overall level of criminal activity, can generate a positive spillover effect that protects all sites. If such spillover effects are not internalized, and other effects are absent, there will be an underincentive for individuals to invest in security.

However, private security goods that are observable to a criminal at the time of a criminal act can simultaneously generate negative spillovers. Specifically, such observable goods can create a diversion effect; that is, they shift the costs of criminal activity to other less protected targets but do not serve as an overall deterrent to criminal and other wealth-transferring activity (Hui-Wen and Png 1994). Thus, the marginal reduction in the probability of an attack faced by a site protected as a result of a marginal increase in security expenditures is not a gross social gain, as it will be partially offset by an increase in the probability, *ceteris paribus*, that other sites will be attacked. One consequence of this diversion effect is that there can be an equilibrium over incentive to invest in observable private security goods.

Moreover, even if the between-site (victim) spillovers mentioned in the preceding paragraph are internalized, private security expenditures can be socially excessive. As noted, when private security expenditures address socially costless transfers, the social objective is to minimize the total resources spent on attempting to achieve such transfers and on preventing such transfers. However, the objective of victims is to minimize the total amount of wealth transferred from them and to minimize the expenditures aimed at preventing such transfers. And the objective of the criminal is to maximize the amount of transfers net of the resources used to achieve the transfers. Because expenditures aimed at reducing the size of the transfers are not socially beneficial, the fact that both the potential criminal and the potential victim take into account the size of the

Table 1.1. *A comparison of observable equilibrium security expenditure levels: private goods case with costless transfers*

	Socially optimal level (x^{**})	Individual level (x^{*})	Cooperative level (x_0)
Socially optimal level (x^{**})		$x_0 > (<)\ x^{**}$ *Ranking between individual and social levels ambiguous*	$x^{*} > x^{**}$ *Cooperatives overinvest*
Individual level (x^{*})			$x_0 > (<)\ x^{*}$ *Ranking between individual and cooperative levels ambiguous*
Cooperative level (x_0)			

transfers in making their individual resource allocations creates a divergence between the social and private incentives to invest in security.

One situation in which the between-site spillovers are internalized is where potential victims agree to collectively set the level of security expenditures. In this situation, the between-site spillovers will be internalized through the agreement setting the collective security level. Thus, to the extent that the individuals' incentives to divert crime to other sites would result in excessively high security expenditures, the collective agreement functions to suppress the individuals' incentives to engage in this type of socially inefficient "arms race." However, the cooperative will take into account the effect that security expenditures will have on the size of the transfers. In contrast, the social calculus ignores this effect on the size of the transfer. As a result, the cooperative will have a greater marginal incentive to invest in security than is socially optimal, and will set a level of expenditures that is above the socially optimal level.

Table 1.1 summarizes the primary results in the case where the criminal activity results in a socially costless transfer, the security goods are observable, and the only social costs are the resources spent by criminals attempting to achieve such transfers and by potential victims attempting to prevent them. In this case, the marginal social benefit from an incremental increase in security expenditures equals the marginal reduction in the resources used by criminals, which equals the decrease in the frequency of an attack times the incremental cost of an attack (Kobayashi forthcoming). If potential victims set security levels

cooperatively, the marginal private benefit will contain this same deterrence effect. A cooperative will also take into account the marginal reduction in the expected magnitude of the transfer, which is a private but not a social benefit. This additional private benefit will cause the cooperative to overinvest in security.

Individuals setting the levels of security noncooperatively may either under- or overinvest in security. The individual's marginal benefit calculation will also take into account, as a private but not a social benefit, the same marginal reduction in the expected magnitude of the transfer taken into account by the cooperative. The individual also takes into account how an incremental expenditure will alter the frequency with which he or she will be attacked. However, this effect is distinct from the reduction in the overall frequency of attacks that yields the marginal social benefit and is part of the cooperative's calculus. Rather, the individual will take into account the reduction in the frequency of attack that results from criminals being diverted from his or her site to others' sites, whether or not any significant overall reduction in the frequency of attacks results. This individual incentive may be larger or smaller than the social deterrent effect. If it is larger, then individuals will set an equilibrium level of security expenditures that will exceed the cooperatively set level and thus the social level. If it is smaller, then individuals will have smaller incentives than cooperatives and may either under- or overspend the social level.

To illustrate the incentives facing agents considering investments in security and to provide a baseline for the discussion in the next section, Figure 1.1 shows the results of a simulation of the individual, social, and cooperative equilibrium levels of security.[7] The model used to generate Figure 1.1 assumes that security expenditures totaling x were produced under constant returns to scale and that the marginal cost of a unit of security equals 1. These security expenditures affect the activity level of criminals by decreasing the gain from criminal wealth-transferring activity. From a social standpoint, the marginal gain in the pure transfer case equals the marginal reduction in the costs of the criminals' efforts. The socially optimal equilibrium level of security x^{**} is reached when the decrease in the marginal cost of the criminals' efforts equals the marginal cost of the additional unit of security incurred by each potential victim. This occurs at the intersection of the social marginal benefit curve and the horizontal line that intersects the vertical axis at 1.

Figure 1.1 also illustrates the cooperative's incentive to overinvest in security. At any level of x, the marginal private benefit to the members of a security

[7] The underlying assumptions used to generate the simulations are described in Kobayashi (forthcoming).

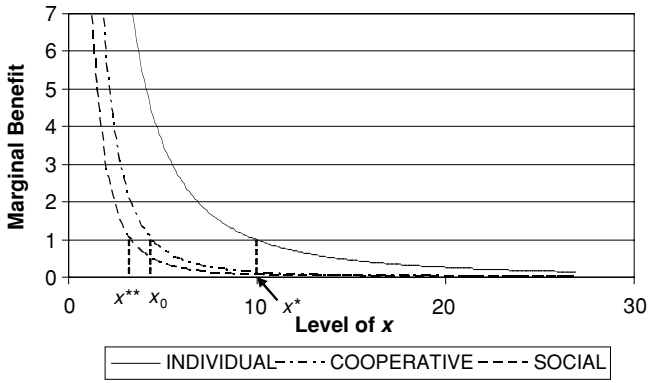

Figure 1.1. Equilibrium security expenditure levels: private goods case with costless transfers.

cooperative equals the social marginal benefit that results from a reduction in the criminals' level of activity plus the private (but not social) benefit associated with the reduction in the size of the transfer. As a result, the cooperative marginal benefit curve lies above the social marginal benefit curve for all levels of x, and thus the cooperative level of security x_0 will be greater than the social level x^{**}.

Figure 1.1 also illustrates a case where the diversion effect results in the individual, noncoordinated level of security (x^*) exceeding both the social (x^{**}) and the cooperative (x_0) levels of security.[8] In order to generate an equilibrium diversion effect, the model assumes that criminals perceive each individual's true level of security x_i with error[9] and will choose to attack the site with the lowest perceived level of protection (Kobayashi forthcoming). Under these conditions, potential victims have a marginal incentive to increase their individual level of security x_i in order to decrease the probability their site will be attacked. As illustrated in Figure 1.1, the incentive to divert criminals to other sites results in an equilibrium level of uncoordinated expenditures that is over three times the socially optimal level of security (Kobayashi forthcoming).

The relative importance of this diversion effect will be dependent upon the technology used to secure individual assets and the ability of criminals to perceive differences in individual security levels. For example, if individual security levels are observed with error, then the importance of the diversion effect

[8] Under the assumptions of the simulation model depicted in Figure 1.1, the social level of security (x^{**}) equals 3.2 units per site, the cooperative level (x_0) equals 4.3 units per site, and the individual, uncoordinated level (x^*) equals 9.9 units per site. For a detailed description of these simulations, see Kobayashi (forthcoming).

[9] Specifically, the criminal observes a proxy variable z_i that equals the actual security level x_i plus a random error term e_i. See Kobayashi (forthcoming).

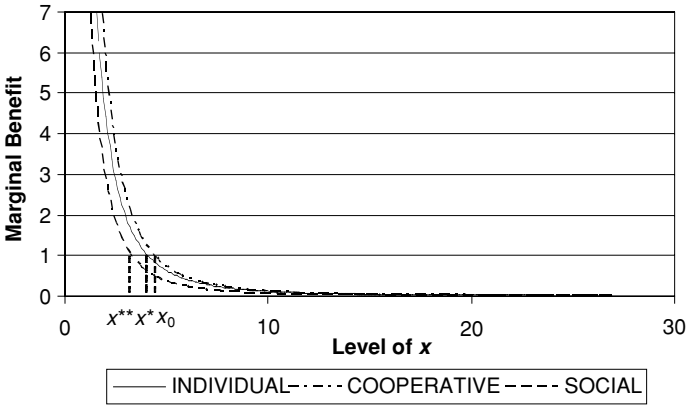

Figure 1.2. Equilibrium security expenditure levels: private goods case with costless transfers and low signal-to-noise ratio.

will depend upon the signal-to-noise ratio of such diversionary expenditures (Kobayashi forthcoming).[10] If the noise level is relatively high, then the diversion effect will be relatively unimportant. However, a relatively low noise level may elevate the magnitude of the diversion effect and create a large individual overincentive to invest in security.

Figure 1.2 shows the result of the simulation when the signal-to-noise ratio is diminished.[11] As shown in the figure, individuals' incentives to expend resources in order to divert attacks to other sites are diminished relative to the case depicted in Figure 1.1. While the simulation depicted in Figure 1.2 results in the individual level of security expenditures (x^*) exceeding the social level (x^{**}), the individual level is below the cooperative level (x_0).[12]

III. PUBLIC AND PRIVATE GOODS

The previous section examined the provision of private security goods such as door locks and security guards. In the cybersecurity context, expenditures on security are likely to be investments in information about the nature and

[10] For a similar analysis of the effect of uncertain legal standards on deterrence, see Craswell and Calfee (1986).

[11] This effect is achieved by assuming that the standard deviation of the random error term e_i is increased by a factor of 10. All other parameters are identical to those used in the simulation that generated Figure 1.1.

[12] Under the assumptions of the simulation model depicted in Figure 1.2, the social level of security (x^{**}) and the cooperative level (x_0) are unchanged. The individual, uncoordinated level (x^*) falls from 9.9 units per site to 4.0 units.

frequency of past attacks, about pending attacks, and about the existence of vulnerabilities to and potential defenses against attacks. Such information is a classic public good that, once produced, can be consumed by multiple sites in a nonrivalrous fashion.[13] Rather than having each site produce its own level of security, efficiency would dictate that these investments in information not be duplicated.[14]

The fact that cybersecurity-related information is a public good alters the analysis in several ways. First, the production of such information is subject to the familiar trade-off between social incentives to allow the free use of already produced information and the incentive to restrict such use to provide incentives for the creation of the information in the first place. Because information produced by one site can be used in a nonrivalrous fashion by other sites, it is not efficient for each site to separately produce its own information. Uncoordinated individual provision of security would likely result in inefficient duplication of effort.

On the other hand, this information cannot be a collective good freely available to all once produced. If security goods are collective goods, then individuals or firms that invest in information and other public security goods will not be able to exclude others from using them, resulting in an incentive to free-ride. An incentive to free-ride acts as a powerful disincentive to produce security, resulting in individual incentives for private security that will be below social levels. Further, the individual incentives to invest in security in order to divert attacks to other sites that cause the overproduction in the private security goods case will not exist in the collective goods case, as other sites would be protected by any collective goods that were produced.

[13] Aviram and Tor (2004) note the nonrivalrous nature of information. In Chapter 8, Tractman notes the public good nature of cybersecurity and the existence of collective action problems.

[14] This does not imply that security goods should be centralized. The analysis and detection of cyberattacks often require an examination of information distributed in a decentralized fashion among many different sites. Thus, a given level of security expenditures distributed over *h* different sites will reduce cybercrime more than the same level restricted to a single site. In other words, the provision of cybersecurity will exhibit network effects (see Chapter 5). Similarly, observations collected from numerous diverse sources may be more valuable than the same number of observations collected from a few firms (Hayek 1945). This analysis suggests that firms have a great incentive to share information in a cybersecurity setting. Similar incentives for sharing of information between competitive firms have raised antitrust concerns. For example, the McCarran Ferguson Act (U.S. Code Title 15, Chapter 20) makes the cooperative gathering of data for the purpose of rate making exempt from the federal antitrust statutes when undertaken by state-regulated insurance companies. For an analysis of information sharing and antitrust issues in the cybersecurity context, see Aviram and Tor (2004). For economic analyses of information sharing between competing firms, see Armantier and Richard (2003), Eisenberg (1981), and Gal-Or (1986).

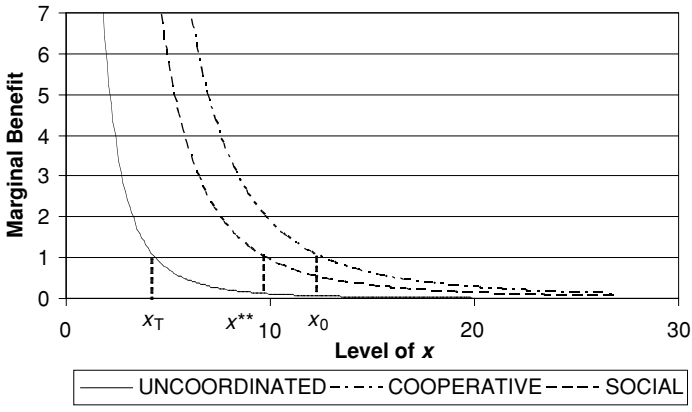

Figure 1.3. Equilibrium security expenditure levels: public goods case with costless transfers.

Figure 1.3 depicts the incentive to invest in security goods that are public in nature. As was the case in the simulations used to generate Figures 1.1 and 1.2, it is assumed that security expenditures totaling x were produced under constant returns to scale and that the marginal cost of a unit of security equals 1. Further, the functional forms for the criminals' cost of effort and gain functions, as well as the number of potential victims and criminals, are identical to those used to generate Figures 1.1 and 1.2.

However, in the simulations used to generate Figure 1.3, each unit of security x can be simultaneously applied to all potential victims. Because each unit of x is not separately incurred by each potential victim, the total level of protection applied to each site at the social optimum is greater than in the private goods case, but the total spending is less. In the private goods case depicted in Figure 1.1, each of sixteen sites spends 3.2 units on security at the social optimum. Thus, the socially optimal total level of security expenditures in the private goods case equals 51.2 units.[15] In the public goods case depicted in Figure 1.3, the socially optimal total level of security expenditures equals 9.7 units, which is applied to all sixteen sites.

In contrast to the private goods case, the uncoordinated level of security expenditures x_T is far below the socially optimal level. As depicted in Figure 1.3, the uncoordinated level of security would equal 4.3 units, compared to the social level of 9.3 units. This level does not equal the per-site expenditure. Rather, it represents an individual site's preference for the total level of expenditures on x by all potential victims. Moreover, while this level of total expenditures satisfies an individual site's first-order conditions, it does not define a unique

[15] See Figure 1.1 and the discussion of this figure in the text.