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