The Economics of Cybersecurity: Principles and Policy Options[1]

Tyler Moore
Center for Research on Computation and Society
Harvard University
Maxwell Dworkin 110
33 Oxford St.
Cambridge, Massachusetts 02138, 02138
tmoore@seas.harvard.edu

Abstract Economics puts the challenges facing cybersecurity into perspective better than a purely technical approach does. Systems often fail because the organizations that defend them do not bear the full costs of failure. For instance, companies operating critical infrastructures have integrated control systems with the Internet to reduce near-term, measurable costs while raising the risk of catastrophic failure, whose losses will be primarily borne by society. As long as anti-virus software is left to individuals to purchase and install, there may be a less than optimal level of protection when infected machines cause trouble for other machines rather than their owners. In order to solve the problems of growing vulnerability and increasing crime, policy and legislation must coherently allocate responsibilities and liabilities so that the parties in a position to fix problems have an incentive to do so. In this paper, we outline the various economic challenges plaguing cybersecurity in greater detail: misaligned incentives, information asymmetries and externalities. We then discuss the regulatory options that are available to overcome these barriers in the cybersecurity context: *ex ante* safety regulation, *ex post* liability, information disclosure, and indirect intermediary liability. Finally, we make several recommendations for policy changes to improve cybersecurity: mitigating malware infections via ISPs by subsidized cleanup, mandatory disclosure of fraud losses and security incidents, mandatory disclosure of control system incidents and intrusions, and aggregating reports of cyber espionage and reporting to the World Trade Organization (WTO).

Keywords: Information security, economics, payment card security, malware, incentives, information asymmetries, externalities, intermediary liability

1 Introduction

Cybersecurity has recently grabbed the attention of policymakers. There have been persistent reports of foreign agents penetrating critical infrastructures, computer compromise facilitating industrial espionage, and faceless hackers emptying thousands of bank accounts. Furthermore, information

[1] This article appears in the *International Journal of Critical Infrastructure Protection,* Volume 3, Issues 3-4, December 2010, Pages 103-117 (http://dx.doi.org/10.1016/j.ijcip.2010.10.002). An earlier version appeared in *Proceedings of a Workshop on Deterring Cyberattacks*, pp. 3—23 © 2010 by the National Academy of Sciences, Courtesy of the National Academies Press, Washington, D.C (http://books.nap.edu/openbook.php?record_id=12997&page=3).

security is now increasingly viewed as a matter of national security. The US military has even recently established Cyber Command to defend the domestic Internet infrastructure and organize military operations in cyberspace.

When considering the national security implications of cybersecurity, it is tempting to think in terms of worst-case scenarios, such as a cyber 'Pearl Harbor' where enemies shut down the power grid, wreak havoc on the financial system, and pose an existential threat. Imagining such worst-case scenarios is useful for concentrating the minds of decision makers and spurring them into action. However, there are downsides to focusing on the most extravagantly conceived threats – it gives the false impression that the situation is so dire that only a radical intervention might help.

In fact, many of the problems plaguing cybersecurity are economic in nature, and modest interventions that align stakeholder incentives and correct market failures can significantly improve a nation's cybersecurity posture. Systems often fail because the organizations that defend them do not bear the full costs of failure. Policy and legislation must carefully allocate responsibilities and liabilities so that the parties in a position to fix problems have an incentive to do so.

In this paper, we outline the key insights offered by an economic perspective on information security, and detail actionable policy recommendations that can substantially improve the state of cybersecurity. In Section 2, we describe four crucial aspects of cybersecurity, for which we later propose policy solutions. First is online identity theft, which is the primary way cyber-criminals steal money from consumers. Second is industrial espionage, where trade secrets are remotely and often undetectably stolen. Third is critical infrastructure protection. The control systems regulating power plants and chemical refineries are vulnerable to cyber attack, yet very little investment has been made to protect against these threats. Finally, we consider botnets, a popular method of attack that impacts nearly all aspects of cybersecurity.

In Section 3, we describe the high-level economic challenges to cybersecurity: misaligned incentives, information asymmetries and externalities. In Section 4, we study how policy may be used to overcome these barriers. We review the different ways liability is assigned in the law, giving an extended discussion to how the law has tackled various Internet vices by exerting pressure on intermediaries, principally Internet service providers (ISPs) and the payment system. Finally, we make four concrete policy recommendations that can improve cybersecurity.

2 Cybersecurity applications

While the intent of this article is to provide generalized advice to help strengthen cybersecurity, it is useful to consider particular applications where cybersecurity is needed. We describe four of the most prescient threats to cybersecurity: online identity theft, industrial cyber espionage, critical infrastructure protection, and botnets.

2.1 Online Identity Theft

One key way in which malicious parties capitalize on Internet insecurity is by committing online identity theft. Banks have made a strong push for customers to adopt online services due to the massive cost savings compared to performing transactions at physical branches. Yet the means of authentication have not kept up. Banks have primarily relied on passwords to identify customers, which miscreants can obtain by simple guessing or by installing 'keystroke loggers' that record the password as it is entered on a computer. Another way to steal passwords takes advantage of the difficulties in authenticating a bank to a consumer. Using a 'phishing' attack, miscreants masquerade as the bank and ask the customer for credentials. Phishing sites are typically advertised via spam email purporting to come from the bank. Keystroke loggers can be installed using a more general ruse – for instance, fraudsters sent targeted emails to the payroll departments of businesses and school districts with fake invoices attached that triggered installation of the malicious software [39].

Once the banking credentials have been obtained, miscreants need a way to convert the stolen credentials to cash. One option is to sell them on the black market: someone who can collect bank card and PIN data or electronic banking passwords can sell them online to anonymous brokers at advertised rates of $0.40–$20.00 per card and $10–$100 per bank account [55]. Brokers in turn sell the credentials to specialist cashiers who steal and then launder the money.

Cashiers typically transfer money from the victim's account to an account controlled by a 'money mule.' The mules are typically duped into accepting stolen money and then forwarding it. The cashiers recruit them via job ads sent in spam e-mails [45] or hosted on websites such as Craigslist or Monster[30], which typically offer the opportunity to work from home as a 'transaction processor' or 'sales executive.' Mules are told they will receive payment for goods sold or services rendered by their employer and that their job is to take a commission and forward the rest, using an irrevocable payment service such as Western Union. After the mule has sent the money, the fraud is discovered and the mule becomes personally liable for the funds.

2.2 Industrial Cyber Espionage

The rise of the information economy has meant that valuable intellectual property of firms is increasingly stored in digital form on corporate networks. This has made it possible for competitors to remotely gain unauthorized access to proprietary information. Such industrial espionage can be difficult to detect, since simply reading the information does not affect its continued use by the victim. Nonetheless, a few detailed cases of espionage have been uncovered. In 2005, 21 executives at several large Israeli companies were arrested for hiring private investigators to install spyware that stole corporate secrets from competitors [57]. In 2009, the hotel operator Starwood sued Hilton, claiming that a Hilton manager electronically copied 100,000 Starwood documents, including market research studies and a design for a new hotel brand [12]. Researchers at the Universities of Toronto and Cambridge uncovered a sophisticated spy ring targeting the Tibetan government in exile [27, 47]. Employees at embassies across the globe were sent emails purporting to be from Tibetan sympathizers. When the employees opened the email attachment, their computers were infected with malware that stole documents and email communications.

Many within government and the defense industrial base argue that industrial cyber espionage is rife. The UK security service MI-5 warned British businesses that Chinese spies were systematically targeting them [32]. The security company Mandiant has claimed that an 'advanced persistent threat' originating in China is being used to systematically steal intellectual property from businesses by infected computers with malware [37]. An anonymous survey of 800 CIOs revealed that many believed they were targeted by espionage, with each firm reportedly losing $4.6 million annually [38]. On the record, however, businesses have remained mum, refusing to acknowledge that the problem poses a significant threat to their profits.

2.3 Critical Infrastructure Protection

It is widely known that the process control systems that control critical infrastructures such as chemical refineries and the power grid are insecure. Why? Protocols for communicating between devices incorporate little, if any, authentication, which potentially allows anyone who can communicate on these networks to be treated as legitimate. Consequently, these systems can be disrupted by a series of crafted messages. The potential for harm was demonstrated by researchers at Idaho National Laboratory who remotely destroyed a large diesel power generator by issuing SCADA commands [42].

In order to carry out an attack, the adversary needs to know quite a bit of specialist knowledge about the obscure protocols used to send the messages, as well as which combination of messages to select. The attacker also needs access to the system. This latter requirement is becoming easier for an attacker to meet due to the trend over the past decade to indirectly connect these control systems to the Internet. The main motivation for doing so is to ease remote administration. A related type of convergence is that the networks themselves are becoming IP-based. That is, the lower level network and transport protocols used to send control messages are now the same as for the wider Internet. This trend also makes it easier for an attacker, once access has been gained, to start sending spurious messages. A few specialist control system engineers understand the transport protocols used by SCADA systems, whereas huge numbers of IT technicians and computer scientists understand Internet protocols. This lowers the technical bar for carrying out attacks.

While many agree that critical infrastructures are vulnerable to cyber attack, few attacks have been realized. Anonymous intelligence officials have reported that Chinese and Russian operatives have regularly intruded into the US electrical grid [23]. Note, however, that no official has gone on the record to describe the intrusions. Nonetheless, the vulnerability cannot be disputed, and the worst case possibility has been demonstrated.

2.4 Botnets

Malware is frequently used to steal passwords and compromise online banking, cloud and corporate services. It is also used to organize infected computers into a 'botnet': a network of thousands or even millions of computers under the control of an attacker that is used to carry out a wide range of services. The services include sending spam, committing online-advertising fraud, launching denial-of-service attacks, hosting phishing attacks, and anonymizing attack traffic. Botnets are different from the previous three categories because they represent an attack method rather than a target. Botnets can be employed in attacks that target all three categories. For instance, some phishing attacks carried out by the rock-phish gang use a botnet infrastructure [43]. The GhostNet/Snooping Dragon espionage of Tibetan authorities utilized a specialized botnet [27, 47]. Finally, botnets are useful for providing anonymous cover for cyber attacks such as those that might harm critical infrastructures.

Botnets are typically crafted for a particular purpose, which vary based on the preferences of the miscreant controlling the botnet, called a 'botnet herder'. Many botnets are designed to simply send spam at the behest of the botnet herder. For example, the Reactor Mailer botnet ran from 2007-2009, at its peak sending more than 180 billion spam messages per day, 60% of the global total [53]; at least 220,000 infected computers participated in the Reactor Mailer botnet each day. The Zeus botnet, by contrast, uses key logger software to steal online credentials which are relayed back to the botnet herder; the botnet is estimated to be as large as 3.6 million computers [60]. Botnets can also be used to carry out denial-of-service attacks. Here, the herder directs the infected computers to make connections to the same website, overloading the targeted site. Botnets were employed to carry out the denial-of-service attacks in Estonia [16] and Georgia [15].

3 Economic Barriers to Improving Cybersecurity

Each of the cybersecurity threats discussed in Section 2 possesses distinct technical characteristics, stakeholders and legal constraints. However, some commonalities remain, notably in the economic barriers inhibiting optimal levels of security investment. We now discuss the crucial common traits first; in Section 4, we will examine the legal and policy options available for each application.

3.1 Misaligned incentives

Information systems are prone to fail when the person or firm responsible for protecting the system is not the one who suffers when it fails. Unfortunately, in many circumstances online risks are allocated poorly. For example, medical records systems are procured by hospital directors and insurance companies, whose interests in account management, cost control, and research are not well aligned with the patients' interests in privacy. Electricity companies have realized substantial efficiency gains by upgrading their control systems to run on the same IP infrastructure as their IT networks.

Unfortunately, these changes in architecture leave systems more vulnerable to failures and attacks, and it is society that suffers most in the event of an outage. Banks encourage consumers and businesses to bank online because it massively reduces branch operating costs, even if the interfaces are not secure and are regularly exploited by attackers. As pointed out by Anderson and Moore [5], misaligned incentives between those responsible for security and those who benefit from protection are rife in IT systems. Consequently, any analysis of cybersecurity should begin with an analysis of stakeholder incentives.

A natural tension exists between efficiency and resilience in the design of IT systems. This is best exemplified by the push over the past decade towards network 'convergence'. Many critical infrastructure systems used to be operated on distinct networks with incompatible protocols and equipment – SS7 protocols managed the phone system, SCADA protocols controlled electrical grids, and so on. It is far cheaper to train and employ engineers whose expertise is in TCP/IP, and run the many disparate applications over a common Internet infrastructure. The downside, however, is that the continued operation of the Internet has now become absolutely essential for each of these previously unconnected sectors, and failure in any one sector can have spillover effects in many sectors. Yet, an individual company's decision to reduce its operating IT costs does not take into account such an increase in long-term vulnerability. Reconciling short-term incentives to reduce operating costs with long-term interest in reducing vulnerability is difficult.

Perfect security is impossible, but even if it were, it would not be desirable. The trade-off between security and efficiency also implies that there exists an optimal level of insecurity, where the benefits of efficient operation outweigh any reductions in risk brought about by additional security measures. For instance, consumers benefit greatly from the efficiency of online banking. The risk of fraud could be reduced to nothing if consumers simply stopped banking online. However, society would actually be worse off because of the added cost of conducting banking offline would outweigh the total losses due to fraud. When misaligned incentives arise, however, the party making the security-efficiency trade-off is not the one who loses out when attacks occur. This naturally leads to suboptimal choices about where to make the trade-off. Unfortunately, such a misalignment is inevitable in many information security decisions.

3.2 Information asymmetries

Many industries report a deluge of data. Some even complain of being overwhelmed. However, in the security space there is a dearth of relevant data needed to drive security investment.

Testifying before the US Congress on March 20, 2009, AT&T's Chief Security Officer Edward Amoroso estimated that the annual profit of cyber criminals exceeds *$1 trillion* [56]. $1 trillion is a lot of money; it is bigger than the entire IT industry, and is approximately 7% of US GDP. It is also likely an extreme overestimate, perhaps triggered by a need to attribute enormous sums to any threat when competing for Congress's attention during this time of trillion-dollar bail-outs.

Note, however, we said it is *likely* an overestimate. The fact is we do not know the true cost of cyber-crime because relevant information is kept secret. Sure, we may never gain access to the miscreants' bank accounts. But we do know that most of revenue-generating cyber-crime is financial in nature, and US banks are not revealing how much they are losing to online fraud.

There is an incentive to under-report incidents across the board. Banks do not want to reveal fraud losses for fear of frightening away customers from online banking; businesses do not want to cooperate with the police on cyber-espionage incidents because their reputation (and their stock price) may take a hit; operators of critical infrastructures do not want to reveal information on outages caused by malicious attack for fear it would draw attention to systemic vulnerabilities. The reticence to share information is only countered by the over-enthusiasm of many in the IT security industry to hype threats.

However, the combination of secrecy and FUD (fear, uncertainty and doubt) is dangerous. To understand why, consider how the used car market works. George Akerlof won a Nobel prize for describing how markets with *asymmetric information*, such as the market for used cars, can fail [2].

Suppose a town has 50 good used cars (worth $2,000 each) for sale, along with 50 ‘lemons’ (worth $1,000 each). The sellers know which type of car they have, but the buyers do not. What will be the market-clearing price? One might initially expect $1,500, but at this price no one with a good car will sell, and so the market price quickly ends up near $1,000. Consequently, the market is flooded with lemons, since no one with a good car would agree to sell at $1,000. The key insight is that buyers are unwilling to pay a premium for quality they cannot measure, which leads to markets with low-quality products.

In 2001, Anderson pointed out that the market for secure software is also a ‘market for lemons’: security vendors may assert their software is secure, but buyers refuse to pay a premium for protection and so vendors become disinclined to invest in security measures [3]. A similar effect is triggered by the refusal to disclose data on losses due to security incidents. The lack of reliable data on the costs of information insecurity make it difficult to manage the risk.

Unreliable information takes many forms, from security vendors overstating losses due to cyber-crime to repeated warnings of digital Armageddon caused by the exploitation of process control system vulnerabilities while suppressing the discussion of realized or attempted attacks. The existence of an information asymmetry does not necessarily mean that society is not investing enough in security, nor that too much money is being allocated. Rather, it simply means that it is likely not investing in the right defenses to the ideal proportion. Ill-informed consumers and businesses are prone to invest in snake-oil solutions if they do not possess an accurate understanding of threats and defenses. Meanwhile, security companies may not be pressured to bring new technologies to market that protect against the most substantial threats. If we do not address the lack of reliable information soon, we are liable to end up with decision makers in industry and government who refuse to implement necessary protections because data that clarifies the magnitude and nature of the most significant threats is just not there.

3.3 Externalities

The IT industry is characterized by many different types of externalities, where individuals’ actions have side effects on others. We discuss three types in turn: network externalities, externalities of insecurity, and interdependent security.

The software industry tends toward dominant firms, thanks in large part to the benefits of interoperability. Economists call this a network externality: a larger network, or a community of software users, is more valuable to each of its members. Selecting an operating system depends not only on its features and performance but also on the number of other people who have already made the same choice. This helps explain the rise and dominance of Windows in operating systems, as well as the dominance of iTunes in online music sales and Facebook in online social networks. Furthermore, it helps explain the typical pattern of security flaws. As a platform vendor is building market dominance, it must appeal to vendors of complementary products as well as to its direct customers. It is more difficult to develop applications for a secure operating system, so security is not emphasized until market dominance has been achieved. Likewise, the opportunities made possible by being first to market explain why insecure software is hurriedly pushed to market, and why software today is issued in perpetual ‘beta’, or test, mode.

Network externalities also help explain why many of the secure upgrades to Internet protocols, such as DNSSEC and S-BGP, have failed to be adopted widely. The security benefits of such protocols are not realized until many other users have also upgraded, which has discouraged early adoption. SSH and IPSec, by contrast, have been much more successful because they provide immediate internal benefits to those who adopt them.

Insecurity creates negative externalities. A compromised computer that has been incorporated in a botnet can pollute the Internet, harming others more than the host. As described in Section 2.4, botnets send spam, host phishing scams, launch denial-of-service attacks, and provide anonymous cover for attackers. In each case, the target of the malicious activity is someone other than the host computer. The societal losses due to control systems failure, such as prolonged power outages, exceed the financial loss to an individual utility in terms of lost revenue. Because the private risks facing utilities are less than the social risks, we would expect an underinvestment in protections against the social risks. Finally, we must also consider the positive externalities of Internet use that go squandered when people are afraid to use the Internet due to its insecurity.

A final type of externality relevant to cybersecurity is interdependent security. Kunreuther and Heal [31] note that security investments can be strategic complements: An individual taking protective measures creates positive externalities for others that in turn may discourage their own investment. Free-riding may result. Varian [59] pointed out that free-riding is likely whenever security depends on the weakest link in the chain: firms do not bother investing in security when they know that other players will not invest,leaving them vulnerable in any case.

4 Prospective Solutions

The economic barriers just discussed – misaligned incentives, information asymmetries and externalities – suggest that regulatory intervention may be necessary to strengthen cybersecurity. We review several different approaches, assess their suitability to the cybersecurity problem, and outline a series of concrete proposals for regulating cybersecurity.

4.1 Overview of Regulatory Options

4.1.1. Ex ante safety regulation vs. Ex post liability

Much of the IT industry has thus far avoided significant regulation. Hence, many of the examples of existing regulatory efforts involving information security concern financial institutions, which face considerably more regulatory scrutiny. *Ex ante* safety regulation is designed to prevent accidents by prescribing safeguards before accidents occur. The bulk of information security regulation (both industry and government led) is compliance-driven, a type of *ex ante* regulation. Firms adopt security policies and 'best practices' and test their own compliance with these rules.

One example of *ex ante* regulation can be found in the Financial Services Modernization Act of 1999 (a.k.a. the Gramm-Leach-Bliley Act), which obliges banks to 'protect the security and confidentiality' of customer information. Federal banking regulators implemented this requirement by specifying processes that banks must comply with, such as adopting a written information security program and establishing programs to assess and manage operational risks. Notably, such regulations avoid technical prescriptions in favor of forcing compliance with organizational requirements. A process-based approach has the advantage of being less dependent on rapidly-changing technologies, as well as making the job of compliance verification easier for regulators. On the other hand, the effectiveness of compliance-driven security policies has been called into question [21]. Given the poor state of cybersecurity, compliance-driven security is at best a qualified failure.

The alternative to proactive *ex ante* regulation is to assign *ex post* liability for failures to the responsible party. Here, the hope is that the threat of monetary damages arising from legal actions will encourage actors to take the necessary precautions to make failures unlikely.

Section 5 of the Federal Trade Commission Act (15 USC § 45) grants the FTC authority to take action against unfair or deceptive acts and practices that affect commerce. Since 2005, the FTC has occasionally charged companies with acting 'unfairly' by failing to adopt reasonable information security practices. Most of their efforts to date have been aimed at non-financial companies that have suffered massive breaches of personal information, including BJs Wholesale Club, DSW and ChoicePoint. Notably, the FTC's awareness of these security failures stems from the proliferation of mandatory breach disclosure regulations adopted by many US states.

Software companies have long avoided any *ex post* liability for vulnerabilities in their own products [7]. Many have argued that making Microsoft liable for the consequences of exploits targeting Windows operating systems would give it a strong incentive to secure its products. This is undoubtedly true, but the question is whether it is too blunt an instrument to incentivize good behavior. For instance, Microsoft has already made huge investments in improving the security of Windows, leading to significant delays in the deployment of Windows Vista. This happened without the threat of liability, although one can argue that it was easier for Microsoft to spend money on security after having established its dominant market position.

A blanket assignment of liability to software developers – say by voiding all contract terms that disclaim liability for defects – is no panacea. First, introducing software liability would create significant negative side effects. The principal negative effect would be a reduction in the pace of innovation. If each new line of code creates a new exposure to a lawsuit, it is inevitable that fewer lines of code will be written. A move towards software liability will also damage the flourishing free software community. Graduate students might hesitate to contribute code to a Linux project if they have to worry about being sued years later if a bug they introduce leads to a critical vulnerability. Resistance to software liability is one of the few points of agreement between open- and closed-source advocates. A second reason is that it is not obvious that introducing liability would make software secure overnight, or even in the long term. This is because software development is inherently buggy. Even responsible software companies that rigorously test for weaknesses do not find all the bugs before a product ships. To expect all software to ship free of vulnerabilities is not realistic.

A better approach, then, is to encourage responsible software development by vendors. Software companies might be required to demonstrate that its software development lifecycle includes adequate testing. The best policy response is to accept that security failures are inevitable, and to instead emphasize robust responses to security incidents (as exemplified by Recommendation 1 in Section 4.2). Furthermore, given the long-standing success of the IT industry in disclaiming software liability, this report focuses on alternative regulatory arrangements more likely to receive broad stakeholder support. *Ex post* liability may still be a viable strategy for other aspects of the cybersecurity, notably process control system security.

Legal scholars have studied the trade-offs between *ex post* liability and *ex ante* regulation regimes. Shavell [52] and Kolstad, Ulen and Johnson [28] find that the best outcome occurs when both are used simultaneously. However, they also find that *ex ante* regulation does not work well when the regulator either lacks information about the possible harm or is uncertain what the minimum standards should be. Unfortunately, both these conditions hold in the context of cybersecurity: security incidents are swept under the rug by affected firms and regulators have yet to find a compliance regime that has significantly improved cybersecurity. Meanwhile, *ex post* liability runs into trouble when firms are not always held liable for harms created or when firms cannot pay full damages. These conditions, too, often hold for cybersecurity. Given this grim reality, we have to consider an alternative approach: information disclosure.

4.1.2 Information disclosure

Since information asymmetries are a fundamental barrier to improving cybersecurity, adopting policies that improve information disclosure may be helpful. Information disclosure has two primary motivations. First is the view, articulated by Louis Brandeis, that 'sunlight is the best disinfectant'. Bringing unfortunate events to light can motivate firms to clean up their acts. Second, disclosure can be motivated by a sense of the community's 'right to know'. The Emergency Planning and Community Right-to-Know Act of 1986 forced manufacturers to disclose to the EPA (and, consequently, the public) the amounts and types of toxic chemicals released into the environment. The aggregated data, known as the Toxic Release Inventory (TRI), has been effective in reducing the amount of toxic chemicals discharged into the environment [29]. The TRI is now available to the public online (www.epa.gov/tri), and citizens can search the database by ZIP code to learn about chemicals and the companies that released them by geographic region. Mandatory information disclosure initiatives such as the TRI are well positioned as a lightweight regulatory alternative to *ex ante* regulation or *ex post* liability.

Another example relevant to cybersecurity is the flurry of privacy breach notification laws adopted in 44 states, led by California in 2002 (California Civil Code 1798.82). Both public and private entities must notify affected individuals when personal data under their control has been acquired by an unauthorized party. The law was intended to ensure that individuals are given the opportunity to protect their interests following data theft, such as when 45 million credit card numbers were stolen from T.J. Maxx's information technology systems [10]. Breach-disclosure laws are also designed to motivate companies to keep personal data secure. Unquestionably, firms are now more aware of the risks of losing personal information, and have directed more investment in preventative measures such as hard drive encryption [46].

Researchers have also found evidence that the information disclosure requirement has both punished violators and reduced harm. Acquisti, Friedman, and Telang [1] found a statistically significant negative impact on stock prices following a reported breach. Meanwhile, Romanosky, Telang, and Acquisti [51] examined identity theft reports obtained from the FTC from 2002 to 2007. Using time differences in the adoption of state breach disclosure laws, they found a small but statistically significant reduction in fraud rates following each state's adoption.

A final benefit of breach-disclosure laws is that they contribute data on security incidents to the public domain. This has reduced the information asymmetry among firms pertaining to the prevalence and severity of leakages of personal information. Unfortunately, there is currently no central clearinghouse for breach reports like the Toxic Release Inventory. Instead, the volunteer website datalossdb.org aggregates reports identified from news reports and letters sent to victims. Despite these limitations, privacy breaches offer the most empirical evidence among all classes of cybersecurity incidents, directly as a result of information-disclosure legislation. For a more complete analysis of the trade-offs between information disclosure, *ex post* liability and *ex ante* regulation, particularly in the context of data breaches, see [50].

However, there are important differences between the circumstances facing toxic chemical and privacy breach disclosures and the types of cybersecurity topics identified in Section 2. One key motivation of existing information disclosure regimes is consumer empowerment. In other words, there is a strong sense of a 'right to know' – notification is required whenever *personal* information is lost, empowering consumers to check credit reports for any suspicious activity. While consumers may also expect to know about cybersecurity incidents, it is often the case that firms lack the requisite information on cyber incidents necessary to invest in countermeasures. If the remote login to the control system of a power station is compromised and the utility keeps mum about what happened, then other power companies will not fully appreciate the likelihood of attack. When banks do not disclose that several business customers have lost millions of dollars due to the compromise of online banking credentials, the customers that have not yet fallen victim are ignorant of the need to take precautions. Thus, in cybersecurity, we face information asymmetries across firms, not just between consumers and firms.

Might information sharing and analysis centers (ISACs) be a viable solution to the asymmetry that exists between firms? ISACs are closed industry groups where participants can voluntarily share security-related information. ISACs were set up by Presidential Decision Directive 63 in 1997 to enable the federal government to coordinate the protection of critical infrastructures (telecommunications, transport, water, chemical plants, banks, etc.), which are primarily owned and operated by private entities.

While ISACs have been useful, they are no substitute for a policy of transparency and information disclosure. Many are classified, so incidents discussed at these meetings are kept secret from many stakeholders as well as the public. The rationale is that companies are more likely to voluntarily participate and be forthright if the information is kept secret. While this may be true, it underscores the value of the mandatory nature of existing information-disclosure efforts described above. (Occasionally, however, particularly egregious incidents are publicized, due to government prodding. For instance, in August 2009 the Financial Services ISAC issued a joint report with FBI the about business-level online banking fraud, describing how criminals had made off with over $100 million, stealing hundreds of thousands of dollars from each victim.) A greater awareness of incidents, including those that companies would rather keep hidden, is made possible by mandatory disclosure. Furthermore, in the cybersecurity domain, competitive interests often preclude voluntary private sector cooperation. For instance, security companies that remove fraudulent phishing websites do not share their data feeds with each other, causing a much slower response [44].

In summary, information disclosure can be a powerful tool for reducing information asymmetries and adjusting misaligned incentives. However, simply righting an information asymmetry will not necessarily fix a problem when externalities are present.

4.1.3 Cyber-Insurance

Insurance is another mechanism for managing the risks presented by network and information security (see, e.g., [9]). A robust market for cyber-insurance would offer several benefits to society. Foremost, insurance could offer a strong incentive for individuals and organizations to take appropriate precautions. Insurance companies could reward security investments by lowering the premiums for less risky actors. Second, because insurance companies base their competitive advantage on risk-adjusted premium differentiation, they have an incentive to collect data on security incidents where claims are made. Consequently, cyber-insurance is often touted as a solution to the informational challenges outlined in Section 3.2. Third, like all types of insurance, cyber-insurance can help firms smooth financial outcomes by accepting the small fixed present cost of an insurance premium in place of future uncertainty of large losses.

Despite these advantages, the market for cyber-insurance has remained small for many years, and has repeatedly fallen short of optimistic growth projections. For instance, a conservative forecast in 2002 predicted the global cyber-insurance market would rise to $2.5 billion by 2005. However, the actual size by 2008 only reached 20% of the forecast for 2005 [6]. Furthermore, the biggest benefits ascribed to cyber-insurance have not been realized. Rather than differentiate premiums by observed security levels, insurance companies base their premiums on non-technical criteria such as firm size. Additionally, insurance companies have not amassed a large claims history that documents incidents.

Why has the market for cyber-insurance been such a disappointment? Factors on both the demand and supply side offer an explanation. On the demand side, insurers complain of a lack of awareness of cyber-risks by firms. In fact, they point to mandatory breach disclosure legislation as a significant step in the right direction, arguing that it has increased awareness at the executive level of this particular category of threat. Consequently, policies that increase the disclosure of cyber risks and incidents would help stimulate further growth in the cyber-insurance market. However, not all demand-side challenges can be addressed by increased awareness alone. The responsibility for dealing with cyber-incidents must be clearly assigned to the appropriate party, otherwise no claims will need to be made. For instance, there is no need for ISPs to take out insurance against computer infections such as viruses and malware when they are not on the hook for mitigation. Legislation that clarifies the liability for cyber incidents would go a long way towards remedying the lack of demand for cyber-insurance.

Barriers to the provision of cyber-insurance extend to issues of supply. First, information asymmetries – in particular, the difficulty of assessing the security of an insured party – can help explain why insurance companies still do not differentiate premiums based on technical criteria. Certification schemes might help, but designing security certifications that cannot be gamed is difficult. Examples of failed certifications include Common Criteria-certified 'tamper-proof' PIN entry devices broken by cleverly-placed paper clips [17] and more malicious websites receiving the TrustE seal of approval than legitimate sites [19]. The other big supply-side problem is that losses from many types of information security risks are globally correlated. Given the dominant market share of the Windows operating system, a new exploit that compromises Windows platforms will affect companies everywhere simultaneously. Whenever such correlations exist, then premiums must be raised, and often the resulting rise in premiums would price many firms out of the market [8]. In practice, insurance companies have avoided such correlations in their claims by adding exclusions to coverage such as excluding damage incurred by untargeted attacks. Such exclusions make cyber-insurance as offered today a far less attractive solution to mitigating risk.

To conclude, cyber-insurance may eventually be part of a long-term solution to improve cybersecurity, but it needs the right mix of policy to help make it viable.

4.1.4 Indirect Intermediary Liability

Perhaps surprising to non-lawyers, liability does not have to be placed on the party directly responsible for harm. Under indirect liability regimes, third parties are held responsible for the wrongs of others. At least three actors are usually involved: the bad actor, the victim, and a third party. A classic example of indirect liability comes from employment law: employers can be held liable for the actions of their employees. Why would indirect liability ever be desirable? Following the logic of Lichtman and Posner [33], a number of conditions can make indirect liability attractive. First, the bad actors could be beyond the reach of the law, either because they cannot be identified or because they could not pay up even if caught. Second, high transaction costs could make designing contracts that assign responsibility infeasible. Once either of these conditions is met, two additional factors should be considered. First, indirect liability is attractive when a third party is in a good position to detect or prevent bad acts. Second, indirect liability is useful when the third party can internalize negative externalities by reducing the incidences of bad acts.

Lichtman and Posner argue that these conditions hold for ISPs in the context of cybersecurity. We defer discussion of the suitability of assigning liability to ISPs for cybersecurity to the next section. For now, we note that while strict liability has been avoided in virtually all Internet contexts, there are some areas where Internet intermediaries have been either obligated or protected from taking actions.

Section 230 of the 1996 Communications Decency Act (CDA) exempted Internet providers from liability for harmful content contributed by its users. Until the CDA was passed, service providers were reluctant to moderate any user posts for fear that doing so would expose them to liability for all the content contributed by users. Section 230 of the CDA offered immunity to service providers that chose to voluntarily delete contributions from users that were deemed inappropriate. Note, however, that the CDA made no *obligation* to remove defamatory or slanderous content, even if it is illegal.

The Digital Millenium Copyright Act (DMCA) of 1998 took a different tack with respect to how service providers respond to online users who violate copyright restrictions. The DMCA also exempts service providers from liability for copyright infringement carried out by their customers. However, there is a catch: ISPs must comply with 'notice-and-takedown' requests from copyright holders by expeditiously removing the content in question in order to obtain the liability exemption.

ISPs are not the only intermediary enlisted by Congress to help rid the Internet of 'bad' actors. Payment networks (i.e., credit card networks such as Visa and MasterCard) are often seen as another intermediary where pressure can be applied. For instance, while early legislation aimed at stopping Internet gambling focused on ISPs, in passing the Unlawful Internet Gambling Enforcement Act (UIGEA) of 2006 Congress ultimately settled on payment processors as the intermediary to assign indirect liability. Payment processors were obliged to implement procedures that stopped Internet gambling transactions. Because Internet gambling operations cannot proceed without credit card payments, leaning on the payment processors was an effective way to shut down operations. Note that the payment system has been used as an intermediary in the fight against other online ills, including child pornography, controlled substances and tobacco sales to minors (see [34] for a thorough explanation for how the law was applied in each case).

Payment card fraud is one area of cybersecurity where indirect liability is already used. The bad actors who commit account fraud victimize cardholders. Under the Truth in Lending Act of 1968 (implemented by the Federal Reserve as Regulation Z), credit card holders are protected from liability for unauthorized charges on their accounts. Similarly, the Electronic Funds Transfer Act (implemented through Regulation E) protects debit card holders from liability for fraudulent use. Instead, the obligation to repay falls on banks that operate the payment system, since the criminals are often out of reach.

It is instructive to examine how liability for payment card fraud has been allocated among intermediaries [35]. In the case of fraud occurring at brick-and-mortar stores, banks rather than merchants traditionally foot the bill. For online transactions, however, the merchant has to pay. This is because online transactions are riskier since the card is not present. Banks and merchants have continued to fight over who should ultimately pay out in different circumstances. The Payment Card System Data Security Standard (PCI DSS) is a series of compliance requirements designed to improve the security of the payment system, particularly for merchants. Merchants found to be non-compliant with PCI requirements are assigned liability for fraud under industry rules. Merchants complain of the high costs of compliance and argue that PCI DSS is nothing more than a thinly veiled, industry-led liability shift from banks to merchants. Banks in turn argue that the issue is fairness, and that merchants must take responsibility for securing payment information and payment systems. A key point when considering what to do about cybersecurity is that any legal ambiguity about which intermediary must pay for remedies is undesirable and can lead to nasty legal battles.

In summary, Congress has acted to regulate illegal online activities by articulating what intermediaries can or must do. There is a range of intervention possible, from 'Good Samaritan' provisions protecting voluntary countermeasures to obligations of action in order to gain exemptions from liability. Most legislative interventions have been hands-off and lightweight, but unafraid to enlist the support of Internet participants to counter undesirable activity.

4.2 Recommendation1: Mitigating malware infections via ISPs by subsidized cleanup

As described in Section 2.4, botnets composed of computers that are infected with malware present a substantial threat to many aspects of cybersecurity. This is because botnets are a preferred tool for carrying out a variety of online attacks. Therefore, in our first recommendation, we describe a way to counter botnets by overcoming the economic barriers described in Section 3 using policies inspired by the regulatory options discussed in Section 4.1.

Recommendation 1: Devise a program of malware remediation with the following attributes:

- ISPs are obliged to act on notifications that their customers are infected with malware by helping to coordinate the cleanup of affected computers. In exchange for cooperation, ISPs receive an exemption from liability for the harm caused by the infected machines. If ISPs do not cooperate, then they become liable for the harm caused by the infected machines.
- The costs of cleanup will be shared between ISPs, government, software companies and consumers.
- Reports of infections (including ISP, platform/operating system, infection vector, time to remediation, remediation technique) must be maintained in a database and made publicly available on the data.gov website.
- Software companies contribute financially to a cleanup fund according to the number of reported infections that affect its software. Software companies receive an exemption from liability for the harm caused by the infected machines in exchange for contributing to the fund.
- Consumer contribution to cleanup is capped at a small fixed amount. Consumers receive guarantees that they will not be disconnected by their ISPs in exchange for cooperating with cleanup efforts.

A substantial portion of Internet-connected computers are infected with malware. Estimates range from a few percent to 25% or more. Malware is frequently used to steal passwords and compromise online banking, cloud and corporate services. It is also used to recruit infected computers into botnets, which may used to send spam, commit online-advertising fraud, launch denial-of-service attacks, host phishing attacks, or anonymize attack traffic.

How does malware get cleaned up today? Sometimes the user will notice. If the user has installed anti-virus software, then the software may detect the malware after receiving updated signatures. However, this often does not work because most malware tries to disable new updates to the anti-virus software. Another option for Windows users is Microsoft's Malicious Software Removal Tool (MSRT). While far from complete, the MSRT automatically detects and removes popular types of malware. If these measures fail, then the user often remains completely ignorant of the presence of malware. However, most malware-infected computers leave a trail of malicious activity that can be identified by third-party security companies that monitor Internet traffic. These companies often notify the relevant ISP of the activity. Some ISPs also actively detect computers that participate in botnets [36] and pass lists of suspected IP addresses to the relevant ISPs. This cooperation stems from ISPs' long-standing cooperation in fighting spam, which is now sent via botnets.

After they are notified about malware on their customers' computers, ISPs have several options. At a bare minimum they can pass along the notice to consumers. In October 2009, Comcast announced a trial program to inform infected customers by browser pop-ups and provide them with instructions for removal [14]. Such notification-only schemes rely on customers to take the necessary steps, which sometimes works for tech-savvy users and malware detectable by tools such as Microsoft's MSRT. Inevitably, though, malware is often not removed by users after they have been notified. To address these cases, Comcast has partnered with McAfee to offer remediation services by skilled technicians for $89.95. Australian ISPs recently announced a notification-based effort [25].

Another ISP-based option is to 'quarantine' infected computers. While in quarantine, users are required to download and install anti-virus software and malware removal tools. They are permitted to rejoin the wider Internet only after the security software is installed and the computer passes a network-based scan for malware. Quarantine is considerably more expensive than notification-only-based interventions because special hardware must be installed at ISPs and more customer-support calls are made. Some ISPs use quarantine, but only for a minority of infected customers. Recently Dutch ISPs announced a program that notifies and quarantines infected customers [20]. Note that in the Dutch and Australian cases many ISPs have joined together in common action. In part, this collective action is designed to allay the fear that customers might switch to a different provider rather than fix the underlying problem.

However, despite the increased interest among some ISPs, by far the most common response by an ISP to a notification about a malware infection is to take no action. Why? The incentive for ISPs to intervene is very weak [58]. Malware harms many victims, from consumers whose credentials are stolen to the targets of DDoS attacks. However, ISPs are not affected much, apart from the prospects of being chided by other ISPs if too many customer machines send out too much spam. By contrast, ISPs face significant tangible costs by intervening. Above all, the costs of customer support in dealing with the phone calls received after sending out notices or placing customers in quarantine are very high. For the ISP, it is much less costly to simply ignore the notifications.

Consequently, the status quo of malware remediation is unacceptable. Many ISPs choose not to act, and those that do avoid cleaning up the hard cases. Notification-only approaches leave many computers infected, while quarantine-based schemes can unfairly shut down the Internet connections of consumers who have followed all the steps but still remain infected. So what should the solution look like?

The first step in a comprehensive solution is to determine who should be responsible for taking action, and how to assign the responsibility. The ISP is a natural candidate for assigning indirect intermediary liability for cleaning up malware. This is because the miscreants who are responsible for the infections are typically beyond the reach of the law. Furthermore, as discussed above, ISPs are in a good position to detect and clean up computers infected with malware. But how should the liability be assigned?

Lichtman and Posner [33] argue for ISPs to take on strict liability for the actions of its customers' computers. In other words, they suggest simply that the ISPs should take the blame for malware-infected customers, and let them choose how they remedy the situation given the threat of legal responsibility. However, considering the exemptions ISPs have historically secured from responsibility for the actions of their customers in other contexts, such an aggressive approach is unlikely to succeed. Instead, we look to the past examples discussed in Section 4.1.4 for inspiration.

The most cautious approach would be to follow the lead of CDA §230 and make cleanup voluntary, explicitly stating that ISPs have no obligation to fix infected computers, but that they are given legal leeway in the event they choose to intervene. While some ISPs are already intervening voluntarily, clarifying the legal right to do so might embolden wary ISPs to act. However, there are distinct disadvantages of this approach. Notably, it does nothing to compensate for the weak incentives that ISPs face in taking action, leading to incomplete remediation. Furthermore, by enshrining a lack of duty, ISPs may choose to intervene even less often than they do in today's more ambiguous environment.

A more ambitious approach (and the one we recommend) is to assign responsibility as has been done in the DMCA. Under a DMCA-like arrangement, ISPs get safe harbor from liability if they clean up infected customer machines upon notification. Notification of infected computers can come from an ISP's own efforts, detection by other ISPs, or from third-party security researchers, as already happens today. Safe harbor is granted if ISPs begin the cleanup process upon notification. They can attempt automated notifications first, and ratchet up their efforts if notifications fail to fix the problems. Quarantine may be tried next, followed by perhaps sending technicians to remediate the machines. Legislation would not be prescriptive in laying out the steps that must be tried and their order; rather, the scheme should be flexible enough to enable ISPs to try different approaches, as long as they are documented and the ultimate solution is a verified, timely cleanup of the affected computer.

ISPs that do not comply with notifications assume liability for the actions of the compromised machines. The amount of liability could be determined by the damages caused. Alternatively, since determining the harm caused by a particular machine is difficult, liability could be assigned as a fixed penalty per ignored infection. Fixed penalties are used in other regulatory contexts. For example, in Europe, airlines are assigned fixed penalties for flight overbooking, cancellations and excessive delays. Fixed penalties are useful because they avoid the problem of quantifying losses following every infringement. The threat of penalties should alter behavior so that, in practice, penalties are rarely issued. Anderson et al. [4] have recommended that the European Commission introduce fixed penalties for ISPs that do not expeditiously comply with notifications of compromised machines in their networks. Such an approach could be effective in our context as well.

Three additional caveats to the designed countermeasure are still needed: a fair distribution of who pays for cleanup, the transparency achieved through mandatory disclosure of reported infections, and consumer protection that ensures Internet connectivity is not threatened by cleanup efforts. We discuss each in turn.

Assigning ISPs the responsibility of ensuring that their infected customers are cleaned up would impose a costly obligation on them. This is somewhat unfair, because it is not an ISP's fault that a user has been infected. But indirect liability regimes need not be fair to be effective. However, a fair allocation of responsibilities can help ensure that the proposal has broad support. Surely, the software companies who designed the insecure systems should bear some responsibility for cleaning up the mess. To that end, we recommend that the costs of cleanup be shared by ISPs, government, software companies and consumers. ISPs already pay as a result of the increased overhead in managing the cleanup process. Governments and software companies should pay by contributing to a fund that will help subsidize the ISP cleanup process. There are already precedents for cost-sharing between third parties in the cybersecurity context. First, Luxembourg is exploring the possibility of subsidizing malware cleanup [13]. Second, as mentioned in Section 4.1.4, banks have negotiated arrangements with merchants to help pay for fraudulent transactions whenever standard security practices have not been met. For instance, Visa negotiated a payment of $40.9 million from TJX to reimburse banks following a breach that affected 46 million cardholders [10], while in January 2010 Heartland agreed to pay MasterCard $41 million following a breach of 100 million credit card numbers [24]. Rather than negotiating one-off settlements between intermediaries, we recommend establishing a fund to receive regular payments from software companies, given the persistent nature of malware infections.

The government should pay for cleanup because it values clean networks and the reduction in denial-of-service attacks, corporate espionage and identity theft made possible by malware. Software companies should pay because holes in their software make the compromises possible. To make participation more palatable, we recommend that, in exchange for helping to pay for the cleanup, software companies be granted safe harbor from any harm the compromised machines have caused prior to cleanup. The payment could be distributed according to what caused the infections. If the infection reports include the method of exploitation, then it is easy to identify if the culprit is Windows XP (Microsoft pays) or Acrobat (Adobe pays). Once the scheme is up and running, the contributions for the succeeding quarter can be based on the share of cleanup costs for the previous quarter. In this way, companies are rewarded for selling software that is more secure. In some cases (e.g., for open source software), it will be difficult to track down the party responsible for developing the software that has been exploited. In this case, the government can pay the unclassified share.

An absolutely critical component of the scheme is that it be transparent. We recommend mandatory disclosure of malware infections and cleanup in the same spirit as the privacy breach notification laws. Rather than requiring companies to notify only consumers of infections, we recommend mandatory disclosure of all de-identified data regarding notification of compromises and the cleanup processes. Reports of infections (including ISP, machine operating system, infection vector, time to remediation, remediation technique) must be reported to a database and made publicly available on the data.gov website. The format for the incident data could adhere to the IODEF standard (xml.coverpages.org/iodef.html).

Mandatory collection and publication of data is an essential component of the scheme and part of the grand bargain between ISPs and software companies that would receive liability exemptions in exchange for cooperation with the cleanup process. Mandatory disclosure of infections will help address the information asymmetry that hampers information security investment. Disclosure would put valuable security incident data in the public domain, and it is likely that it will trigger the same 'sunshine effect' as in the cases of environmental pollution due to the Toxic Release Index and personal information protection as a result of breach-disclosure laws. Some of the worst offenders (both ISPs and software companies) will be uncovered, raising awareness of the problem and providing an incentive for investment in defense. Progress will become measurable, not only to insiders but also to outsiders on the scale and quality of cleanup efforts. Public disclosure will help companies gain trust in the level of financial contributions required for assisting cleanup. Finally, transparent disclosure helps give credibility to the claim that improving cybersecurity is taken seriously at the government level. If the US can demonstrate its commitment to cleaning up its own networks, then the resulting improvements in security can be used to apply pressure on other countries to follow suit.

We have already staked out the roles for governments, ISPs and software companies. What of consumer responsibility? Even customers who adhere to all the best practices may become infected. According to Panda Security (www.pandasecurity.com/img/enc/infection.htm), 3.94% of US computers scanned were actively running high-risk malware at the time of the scan; 8.21% of computers without antivirus software were running high-risk malware, but so did 1.64% of computers *with* antivirus software. Furthermore, attackers may craft 'zero-day' exploits – attacks that exploit vulnerabilities previously unknown to the software provider or antivirus company – that no software can defend against. Finally, contrary to popular belief, getting infected is not caused by 'irresponsible' web browsing habits such as visiting disreputable websites and installing dubious programs willy-nilly. A common method of compromise is the 'drive-by-download', where miscreants compromise popular websites so that when unsuspecting users visit the website, the site secretly downloads and installs malware onto the computer. In one study, researchers at Google found 3 million drive-by-download URLs, and that 1.3% of Google's incoming search queries return at least one drive-by-download link in its results [49].

Taken together, the evidence points to a situation where users cannot easily be blamed when malware takes over their computer. But in an economic analysis of liability, fairness takes a back seat to identifying the party in the best position to efficiently address the problem. Consumers are generally not in a good position to defend themselves. They do not write the buggy software, and so they cannot plug the holes; they do not have a network-level view of Internet traffic, so they cannot determine whether or not they are infected (as ISPs can). At best, they can take some safety precautions such as patching their computers and installing antivirus software. There is little more we can expect from consumers, and even if all the consumers were to automatically install patches and run antivirus software, the problem would remain. Consequently, consumers are not in the best position to fix the problem.

In light of this reality, policy should focus on ensuring that consumers are protected in the course of any cleanup efforts. Consequently, we recommend that the financial responsibility placed on the user be limited. Again, we have a precedent from the financial industry in Regulations E and Z, where payment card holders are not liable for fraudulent activity beyond a small fixed amount. A small remediation fee, capped at around $20 or so, would make the cleanup process smoother for malware victims while minimizing the moral hazard for users. Perhaps, the fee could be slightly higher for users who do not have antivirus software installed.

It is also essential that the burden on the ISP is to actually remedy the infection. Disconnecting the Internet connections of users is not an acceptable option, given the increasing reliance on the Internet for basic services. The only exception allowing disconnection could be if consumers do not cooperate with ISP cleanup efforts. Otherwise, ISPs should have a duty to perform cleanup, thereby capping the out-of-pocket expenses for consumers. This addresses the concern that ISPs might find it cheaper to kick off less profitable subscribers rather than clean up their machines.

4.3 Recommendation 2: Mandated disclosure of fraud losses and security incidents

Our second recommendation is considerably simpler than the first.

Recommendation 2: Establish a program that regularly publishes the aggregated loss figures related to online banking and payment cards on data.gov.

The aggregated loss figures would include the following:

- *Incident figures*: number of incidents, total value stolen, total value recovered for specified number of incidents
- *Victim bank demographics*: number of banks affected, number of customer accounts impacted per bank, monetary loss per customer, bank type, precautions taken by bank (e.g., two-factor authentication, back-end controls used)
- *Victim customer demographics*: separate tally of incidents and monetary losses for business and retail customers
- *Attack vector (where known)*: number of incidents and monetary losses for each attack vector (e.g., keystroke-logging malware, phishing, credit card skimming, payment network compromise)
- *Business category*: number of incidents and monetary losses for the categories of online banking, payment cards (transaction type: retail, card present, card not present), and ATM fraud

At present, no objective measures exist to answer the seemingly straightforward questions: is online identity theft increasing or decreasing? How many people and businesses fall victim to fraud online, and how much money is lost? Are online banking and e-commerce less safe than transactions in the real world? Without a way to answer these questions, effective policy cannot be developed to improve cybersecurity.

Fortunately, a low-cost solution is readily available: require financial institutions to report back on fraud losses and aggregate their responses. It is not as if this information has to be kept secret. Banks in Spain, Britain and Australia regularly disclose aggregate information on payment card fraud. In 2009, for example, UK banks lost £440 million (approximately $641 million) due to all forms of payment fraud, while £59.7 million ($87 million) was attributed to online banking in particular [48]. Richard Sullivan, an economist at the Federal Reserve, has argued that fraud statistics should be published in order to get a better grip on fraud levels and provide information on whether investments to secure the payment card infrastructure are needed [53].

Within the US, there are some existing efforts to collect data on online fraud. David Nelson at FDIC has attempted to collect fraud figures from US banks on a voluntary basis. He estimates that $120 million was collectively lost by US banks due to malware infections targeting online banking services [40]. The FBI runs the Internet Crime Complaint Center (IC3), which invites members of the public to submit reports on a wide variety of Internet scams. Some aggregate figures are made available by the IC3 in annual reports (see, e.g., [26]), but access to most of the IC3 data is available only to law enforcement. The Financial Crimes Reporting Center collects suspicious activity reports from banks, but these mainly focus on money laundering activity. The Financial Services ISAC shares confidential, high-level information on threats between banks.

These efforts exhibit a number of significant limitations compared to the mandatory disclosure scheme we recommend. First, the reports are voluntary in nature, making them incomplete, unrepresentative of limited use to draw reliable trends. Very few privacy breaches were disclosed until the California law was passed, and we might suspect that the reports of online fraud are inaccurate estimates of reality. In the case of IC3, the trouble is that quantifying losses is difficult in many circumstances, primarily because it relies on self-reporting. Second, the detailed reports are often secret – IC3 reports are shared only within law enforcement, the FS-ISAC is closed, and so on. Finally, efforts such as the FDIC tally of fraud figures are one-off samples, which make inferring trends over time impossible.

The principal justification for *mandating* the public disclosure of incidents and losses is that the financial industry does not internalize all the costs of insecurity. Consumers are protected by Regulations E&Z, but businesses are not, and merchants are expected to help cover the costs of fraud. If banks instead choose to cover all losses, then publishing loss figures is less crucial. As it stands, banks do not internalize all costs, and so the public deserves a fair and transparent accounting of the share paid by each entity. This is why it is recommended to disclose, in addition to aggregated loss figures, a breakdown of the number and average loss of incidents for consumers and businesses. Additionally, it is important to know the distribution of losses between banks and merchants. These types of information can help answer questions such as how many people's lives are being disrupted by online fraud, if any groups pay a disproportionate share, and if this changes over time.

A second motivation for mandated disclosure is that payment systems exhibit significant network externalities. Visa, Mastercard and American Express have cultivated a very successful credit card network with millions of participating merchants and cardholders. The value of this user base is enormous, and presents a significant barrier to would-be new entrants who might offer more secure payment alternatives. Having already invested heavily in a less secure payment technology and achieved market dominance, existing payment networks may be reluctant to invest further in security mechanisms to reduce fraud that is borne in part by third parties. Payment networks might argue that they are already investing in security, and point to the efforts undertaken in Europe to upgrade to PIN-based smartcard authentication.

Credible reporting of financial fraud losses can settle disputes over whether enough is being done, and it can serve as useful motivation for funding improvements to the security of the financial infrastructure. For instance, banks and payment operators are weighing whether to upgrade the payment network infrastructure to a more secure smartcard-based system [35]. Comprehensive fraud statistics would help banks *and* merchants determine if there has been a substantial increase in card-not-present fraud to justify further security investments. Similarly, the National Strategy for Trusted Identities in Cyberspace being pitched by the White House needs private sector buy-in to be successful, but this will not happen unless firms believe that improvements to online authentication are needed. How can firms agree to spend on security when they do not have an accurate picture of how much is being lost due to presence of a less secure infrastructure? Publishing regular statistics on losses now will motivate future investments if the problem is truly as big as it is often claimed.

4.4 Recommendation 3: Mandated disclosure of control system incidents and intrusions

Anonymous intelligence officials have released stark warnings that Chinese and Russian operatives have regularly intruded into the US electrical grid [23]. In June 2010 a Belarusian security firm uncovered Stuxnet, a worm possibly targeting Iranian control systems made by Siemens [41]. Apart from Stuxnet, no documented case of a successful cyber attack on process control systems has been publicly presented. (It is even unclear whether Stuxnet should be viewed as 'successful' attack, since no physical damage has been found to have been caused by the worm.) In fact, when researchers from the Tuck School of Business interviewed an oil and gas refiner as part of a field study [18], they were told by the VP for refining that he 'had never heard of' a cyber incident shutting down a plant in the industry. The VP went on to state that he would consider investing in process control systems security only after a similar-sized refinery was attacked.

Such different perspectives are hard to reconcile – attacks are supposed to be pervasive, but operators on the ground have yet to observe a single incident. One possible explanation is that the reports of incidents are exaggerated. Many of the individuals who are sounding the alarm stand to gain from increased security investments. Alternatively, the existing mechanisms for exchanging information, the sector-specific ISACs, have failed. ISACs have been in operation for around a decade, which is sufficient time to assess the effectiveness of the voluntary, closed-door information exchanges. Either ISACs have failed to effectively communicate the severity of threats to relevant parties in industry, or there has not been much to report.

Fortunately, there is a reasonable way to get to the bottom of this conundrum: require mandatory disclosure of all cyber incidents and intrusions to regulators, with a substantial public reporting capacity. If the intrusions are in fact happening, the entities who detect the intrusions should have a duty to report them. In fact, the ISACs could serve as the organizations that receive the reports, provided that there is a clear duty on the part of the ISACs to produce public reports that are widely disseminated. **Recommendation 3: Mandatory disclosure of control system incidents and intrusions to the relevant ISACs, who are responsible for further public dissemination.**

There has been some tentative movement in this direction within the electricity industry. The self-regulatory body NERC requires power companies to report to regulators whenever they observe a disturbance suspected to have been caused by sabotage (NERC standard CIP-001). The reports themselves are kept secret, and as far as we know, are not shared with other firms in the industry. This is a useful start because it demonstrates an interest in keeping track of malicious disruptions. However, it is limited in the sense that reporting is only required when an outage occurs. Detecting that Chinese spies have penetrated the administrative interface of a SCADA system does not have to be reported, unless it caused the power to go out. Also, there is no explicit requirement to share the reported information with other utilities, which does not help the oil refiner who is waiting to invest in security until he hears about another industry entity being attacked.

It must be mentioned that mandatory disclosure is no panacea. Disclosure will help address the lack of information on incidents, but the long-tailed nature of cyber attacks on process control systems means that the effort could yield few reports. Furthermore, the problem of externalities persists.

4.5 Recommendation 4: Aggregate reports of cyber espionage and provide reports to the WTO

Industrial espionage is claimed to be a significant problem for US companies. However, the affected companies are naturally reticent to publicly discuss their experiences out of fear that their stock price may take hits. Perhaps, though, the thinking is starting to change. In January 2010, Google disclosed that it had been the victim of a cyber attack focused on industrial espionage that apparently originated in China [11]. Subsequently it was revealed that at least 34 companies were affected, including Yahoo, Symantec, Northrop Grunman and Dow Chemical.

Unfortunately, since the trade secrets were believed to be stolen by an international entity, the US Uniform Trade Secrets Act and Economic Espionage Act cannot easily be enforced. This does, however, leave one option: the TRIPS agreement of the WTO. Deciding to bring cases to the WTO is always politically delicate. However, if the US suspects that industrial espionage is rife, and largely coming from a single country (i.e., China), then it may be worth the effort to prepare a WTO complaint. It is true that such a complaint could negatively affect the stock prices of the firms that are named as victims. Nevertheless, if espionage is anywhere near as pervasive as what has been uncovered in the Google case, then it may be in the strategic interest of the US to take action.

5 Conclusions

An economic perspective is essential to understand the state of cybersecurity today, as well as how to improve it moving forward. In this paper, we have described several key economic challenges: misaligned incentives, information asymmetries and externalities. We have also reviewed the policy options available for overcoming these barriers, notably information disclosure and intermediary liability. Our principal recommendations are to encourage ISPs to take a more active role in cleaning up infected computers, and to collect and publish data on a range of security incidents. These recommendations are designed to raise awareness of cybersecurity issues and assign responsibility for action by the private sector so that the risks to society may be mitigated.

References

1. A. Acquisti, A. Friedman and R. Telang, Is there a cost to privacy breaches? An event study, *Proceedings of the Twenty-Seventh International Conference on Information Systems* (citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.73.2942&rep=rep1&type=pdf), 2006.

2. G. Akerlof, The market for "lemons": Quality uncertainty and the market mechanism, *The Quarterly Journal of Economics*, vol. 84(3), pp. 488-500, 1970.

3. R. Anderson, Why information security is hard: An economic perspective, *Proceedings of the Seventeenth Annual Computer Security Applications Conference,* pp. 358-365, 2001.

4. R. Anderson, R. Böhme, R. Clayton and T. Moore, Security economics and European policy, in *Managing Information Risk and the Economics of Security*, M. Johnson (Ed.), Springer, New York, pp. 55–80, 2008.

5. R. Anderson and T. Moore, The economics of Information Security, *Science*, vol. 314(5799), pp. 610-613, 2006.

6. T. Bandyopadhyay, V. Mookerjee and R. Rao, Why IT managers do not go for cyber-insurance products, *Communications of the ACM,* vol. 52(11), pp. 68--73, 2009.

7. D. Barnes, Deworming the Internet, *Texas Law Review,* vol. 83(1), pp. 279-329, 2004.

8. R. Böhme and G. Kataria, Models and measures for correlation in cyber-insurance, *Proceedings of the Fifth Workshop on the Economics of Information Security* (weis2006.econinfosec.org/docs/16.pdf), 2006.

9. R. Böhme and G. Schwarz, Modeling cyber-insurance: Towards a unifying framework, *Proceedings of the Ninth Workshop on the Economics of Information Security* (weis2010.econinfosec.org/papers/session5/weis2010_boehme.pdf), 2010.

10. *Business Wire*, Visa and TJX agree to provide U.S. issuers up to $40.9 million for data breach claims (www.businesswire.com/portal/site/home/permalink/?ndmViewId=news_view&newsId=20071130005355), November 30, 2007.

11. A. Cha and E. Nakashima, Google China cyberattack part of vast espionage campaign, experts say, *The Washington Post*, January 14, 2010.

12. A. Clark, Starwood sues Hilton for "stealing trade secrets", *The Guardian*, April 17, 2009.

13. R. Clayton, Might governments clean-up malware? *Proceedings of the Ninth Workshop on the Economics of Information Security (*weis2010.econinfosec.org/papers/session4/weis2010_clayton.pdf), 2010.

14. Comcast, Comcast Unveils Comprehensive “Constant Guard” Internet Security Program (www.comcast.com/About/PressRelease/PressReleaseDetail.ashx?prid=926), October 8, 2009.

15. D. Danchev, Coordinated Russia vs. Georgia cyber attack in progress, *ZDNet* (www.zdnet.com/blog/security/coordinated-russia-vs-georgia-cyber-attack-in-progress/1670), August 11, 2008.

16. J. Davis, Hackers take down the most wired country in Europe, *Wired Magazine*, vol. 15(9) (www.wired.com/politics/security/magazine/15-09/ff_estonia?currentPage=all), August 21, 2007.

17. S. Drimer, S. Murdoch and R. Anderson, Thinking inside the box: System-level failures of tamper proofing, *Proceedings of the IEEE Symposium on Security and Privacy*, pp. 281-295, 2008.

18. S. Dynes, E. Goetz and M. Freeman, Cybersecurity: Are economic incentives adequate? in *Critical Infrastructure Protection*, E. Goetz and S. Shenoi (Eds.), Springer, Boston, Massachusetts, pp. 15-27, 2007.

19. B. Edelman, Adverse selection in online “trust” certifications, *Proceedings of the Eleventh International Conference on Electronic Commerce*, pp. 205-212, 2009.

20. G. Evron, Gadi, Dutch ISPs sign anti-botnet treaty, Dark Reading (www.darkreading.com/blog/archives/2009/09/dutch_isps_sign.html), September 29, 2009.

21. Forrester Consulting, The Value of Corporate Secrets: How Compliance and Collaboration Affect Enterprise Perceptions of Risk, (www.rsa.com/products/DLP/ar/10844_5415_The_Value_of_Corporate_Secrets.pdf), March, 2010.

22. N. Gohring, Heartland, MasterCard settled over data breach, *PCWorld* (www.pcworld.com/businesscenter/article/196711/heartland_mastercard_settle_over_data_breach.html), May 19, 2010.

23. S. Gorman, Electricity grid in U.S. penetrated by spies, *The Wall Street Journal*, April 8, 2009.

24. K. Higgins, “Aurora” attacks still under way, investigators closing in on malware creators, *Dark Reading* (www.darkreading.com/vulnerability_management/security/attacks/showArticle.jhtml?articleID=222700786), February 10, 2010.

25. J. Hilvert, eSecurity code to protect Australians online, Internet Industry Association (iia.net.au/index.php/section-blog/90-esecurity-code-for-isps/757-esecurity-code-to-protect-australians-online.html), September 11, 2009.

26. IC3, 2009 Internet Crime Report (www.ic3.gov/media/annualreport/2009_IC3Report.pdf), 2009.

27. Information War Monitor, Tracking GhostNet: Investigating a cyber espionage network (www.scribd.com/doc/13731776/Tracking-GhostNet-Investigating-a-Cyber-Espionage-Network), 2009.

28. C. Kolstad, T. Ulen and G. Johnson, Ex Post liability for harm vs. Ex Ante safety regulation: Substitutes or complements? *American Economic Review,* vol. 80(4), pp. 888-901, 1990.

29. S. Konar and M. Cohen, Information as regulation: The effect of community right to know laws on toxic emissions, *Journal of Environmental Economics and Management,* vol. 32(1), pp. 109-124, 1997.

30. B. Krebs, “Money Mules” help haul cyber criminals’ loot, *The Washington Post*, January 25, 2008.

31. H. Kunreuther and G. Heal, Interdependent security, *Journal of Risk and* Uncertainty, vol. 26(2—3), pp. 231—249. 2003.

32. Leppard, China bugs and burgles Britain, *The Sunday Times*, January 31, 2010.

33. Lichtman and E. Posner, Holding internet service providers accountable, in *The Law and Economics of Cybersecurity*, M. Grady and F. Parisi (Eds.), Cambridge University Press, New York, pp. 221-258, 2006.

34. M. MacCarthy, What Internet intermediaries are doing about liability and why it matters, *ExpressO* (works.bepress.com/mark_maccarthy/1), 2009.

35. M. MacCarthy, Information security policy in the U.S. retail payments industry, *Proceedings of the Ninth Workshop on the Economics of Information Security* (weis2010.econinfosec.org/papers/panel/weis2010_maccarthy.pdf), 2010.

36. MAAWG, Mitigating large-scale bot infections in residential networks (www.maawg.org/sites/maawg/files/news/MAAWG_Bot_Mitigation_BP_2009-07.pdf), July 2009.

37. Mandiant, Mandiant releases inaugural M-trends Report at US Department of Defense Cyber Crime Conference (www.mandiant.com/news_events/article/mandiant_releases_first_annual_m-trends_report_at_u.s._department_of_d/), January 27, 2010.

38. McAfee, Unsecured economies: Protecting vital information (www.cerias.purdue.edu/assets/pdf/mfe_unsec_econ_pr_rpt_fnl_online_012109.pdf), 2009.

39. L. McGlasson, FDIC warns of online fraud against banks, small businesses, *BankInfoSecurity* (www.bankinfosecurity.com/articles.php?art_id=1732), August 26, 2009.

40. R. McMillan, FDIC: Hackers took more than $120M in three months, *Computerworld* (www.computerworld.com/s/article/9167598/FDIC_Hackers_took_more_than_120M_in_three_months?source=rss_news), March 8, 2010.

41. R. McMillan, Siemens: Stuxnet worm hit industrial systems, *Computerworld*, (http://www.computerworld.com/s/article/print/9185419/Siemens_Stuxnet_worm_hit_industrial_systems?taxonomyName=Network+Security&taxonomyId=142), September 14, 2010.

42. J. Meserve, Mouse click could plunge city into darkness, experts say, *CNN* (www.cnn.com/2007/US/09/27/power.at.risk/index.html), September 27, 2007

43. T. Moore and R. Clayton, Examining the impact of website take-down on phishing, *Proceedings of the Anti-Phishing Working Group eCrime Researchers Summit*, pp. 1–13, 2007.

44. T. Moore and R. Clayton, The consequence of non-cooperation in the fight against phishing, *Proceedings of the Anti-Phishing Working Group eCrime Researchers Summit*, pp. 1–14, 2008.

45. T. Moore and R. Clayton, The impact of incentives on notice and take-down, in *Managing Information Risk and the Economics of Security*, M. Johnson (Ed.), Springer, New York, pp. 199–223, 2008.

46. Mulligan and K. Bamberger, Security Breach Notification Laws: Views from Chief Security Officers, Samuelson Law, Technology and Public Policy Clinic, University of California-Berkeley School of Law, Berkeley, California (www.law.berkeley.edu/files/cso_study.pdf), 2007.

47. S. Nagaraja and R. Anderson, The Snooping Dragon: Social-malware Surveillance of the Tibetan Movement, Technical Report UCAM-CL-TR-746, University of Cambridge Computer Laboratory, Cambridge, United Kingdom (www.cl.cam.ac.uk/techreports/UCAM-CL-TR-746.pdf), 2009.

48. Payments News, UK card and banking fraud losses down 28% in 2009 to £440.3M, (www.paymentsnews.com/2010/03/uk-card-and-banking-fraud-losses-down-28-in-2009-to-4403mm.html), March 2010.

49. N. Provos, P. Mavrommatis, M. Rajab and F. Monrose, All your iframes point to us, *Proceedings of the USENIX Security Symposium*, pp. 1-15, 2008.

50. S. Romanosky and A. Acquisti, Privacy costs and personal data protection: economic and legal perspectives of ex ante regulation, ex post liability and information disclosure. *Berkeley Technology Law Journa*l 24(3), 2009.

51. S. Romanosky, R. Telang and A. Acquisti, Do data breach disclosure laws reduce identity theft? *Proceedings of the Seventh Workshop on the Economics of Information Security* (ssrn.com/paper=1268926), 2008.

52. S. Shavell, A model of the optimal use of liability and safety regulation, *RAND Journal of Economics*, vol. 15(2), pp. 271-280, 1984.

53. H. Stern, The rise and fall of reactor mailer, *Proceedings of the MIT Spam Conference* (projects.csail.mit.edu/spamconf/SC2009/Henry_Stern), 2009.

54. R. Sullivan, The Benefits of Collecting and Reporting Payment Fraud Statistics in the United States, Payment Systems Briefing, Federal Reserve Bank of Kansas City, 2009.

55. Symantec, Symantec Global Internet Security Threat Report (eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_internet_security_threat_report_xiii_04-2008.en-us.pdf), 2008.

56. United States Senate, Statement of Edward Amoroso Before the United States Senate Committee on Commerce, Science and Transportation, Hearing on Improving Cybersecurity (commerce.senate.gov/public/?a=Files.Serve&File_id=e8d018c6-bf5f-4ea6-9ecc-a990c4b954c4), March 19, 2009.

57. D. Urquhart, London couple remanded in Israel's biggest industrial espionage case, *The Guardian*, May 31, 2005.

58. M. van Eeten and J. Bauer, The Economics of Malware: Security Decisions, Incentives and Externalities, OECD Science, Technology and Industry Working Paper No. 2008/1, 2008.

59. H. Varian, System reliability and free riding, in *Economics of Information Security*, *Vol. 12, Advances in Information Security*, L. J. Camp, S. Lewis, (Eds.), Kluwer Academic Publishers, Boston, Massachusetts, pp. 1–15, 2004.

60. S. Vaughan-Nichols, Big botnets and how to stop them, *Computerworld* (www.computerworld.com/s/article/9177574/Big_botnets_and_how_to_stop_them), June 2, 2010.