Introducing the Economics of Cybersecurity: Principles and Policy Options

Tyler Moore
Harvard University

ABSTRACT

The economics of information security has recently become a thriving and fast-moving discipline. 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. So 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 in greater detail the various challenges plaguing cybersecurity: misaligned incentives, information asymmetries and externalities. We then discuss the regulatory options that are available to overcome these economic 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 WTO.

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 security is now increasingly viewed as a matter of national security. The U.S. 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 our enemies shut down the power grid, wreak havoc on our 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 our 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 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 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 which we later propose policy solutions for. 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 impacting 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 now 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 customer's bank and ask 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.[1]

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

[1] http://www.bankinfosecurity.com/articles.php?art_id=1732.

rates of $0.40–$20.00 per card and $10–$100 per bank account.[2] 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 (Moore and Clayton 2008a) or hosted on websites such as Craigslist or Monster,[3] which typically offer the opportunity to work from home as a "transaction processor" or "sales executive." Mules are told they will receive payments 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 already sent.

2.2 Industrial Cyber Espionage

The rise of the information economy has meant that the valuable property of firms is increasingly stored in digital form on corporate networks. This has made it easier 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.[4] 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.[5] Researchers at the Universities of Toronto and Cambridge uncovered a sophisticated spy ring targeting the Tibetan government in exile (Information War Monitor 2009, Nagaraja and Anderson 2009). 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 communications.

Many within government and the defense industrial base argue that, rather than occurring in a few isolated incidents, industrial cyber espionage is rife. The UK security service MI-5 warned British businesses that Chinese spies were systematically targeting them.[6] 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 computers infected with malware.[7] An anonymous survey of 800 CIOs revealed that many believed they were targeted by espionage, with each firm reportedly losing $4.6 million annually.[8] On the record, however, businesses have remained mum, refusing to acknowledge the problem as such 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 do not include any authentication, which means that anyone that can communicate on these networks is treated as legitimate. Consequently, these systems can be disrupted by receiving a series of crafted

[2]http://eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_internet_security_threat_report_xiii_04-2008.en-us.pdf.
[3]http://www.washingtonpost.com/wp-dyn/content/story/2008/01/25/ST2008012501460.html.
[4]http://www.guardian.co.uk/world/2005/may/31/israel.
[5]http://www.guardian.co.uk/business/2009/apr/17/industrial-espionage-hotel-industry-lawsuit.
[6]http://www.timesonline.co.uk/tol/news/uk/crime/article7009749.ece.
[7]http://www.mandiant.com/news_events/article/mandiant_releases_first_annual_m-trends_report_at_US_department_of_d/.
[8]http://www.cerias.purdue.edu/assets/pdf/mfe_unsec_econ_pr_rpt_fnl_online_012109.pdf.

messages. The potential for harm was demonstrated by researchers at Idaho National Laboratory who remotely destroyed a large diesel power generator by simply issuing SCADA commands.[9]

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. She 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. Only a few control system engineers understand the transport protocols used by SCADA systems, whereas huge numbers of IT technicians and computer scientists understand Internet protocols. This certainly 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 have regularly intruded into the U.S. electrical grid.[10] 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 place 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 targeting all three of the above categories. For instance, some phishing attacks carried out by the rock-phish gang use a botnet infrastructure (Moore and Clayton 2007). The GhostNet/Snooping Dragon espionage of Tibetan authorities utilized a specialized botnet. 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 (Stern 2009). At least 220,000 infected computers participated in the Reactor Mailer botnet each day. The Zeus botnet, by contrast, includes key logger software to steal online credentials which are relayed back to the botnet herder, and is estimated to be as large as 3.6 million computers.[11] Botnets can also be used to carry out denial-of-service attacks. Here, the herder directs the bots to make connections to the same websites, overloading the targeted site. Botnets were employed to carry out the denial-of-service attacks in Estonia[12] and Georgia.[13]

While the size of botnets varies, the more important factor is what purpose they are being put toward. The Conficker botnet was huge, infecting millions of computers,[14] but has not been associated

[9]http://www.cnn.com/2007/US/09/27/power.at.risk/index.html.
[10]http://online.wsj.com/article/SB123914805204099085.html.
[11]http://www.computerworld.com/s/article/9177574/Big_botnets_and_how_to_stop_them.
[12]http://www.wired.com/politics/security/magazine/15-09/ff_estonia?currentPage=all.
[13]http://www.zdnet.com/blog/security/coordinated-russia-vs-georgia-cyber-attack-in-progress/1670.
[14]http://news.techworld.com/security/114307/experts-bicker-over-conficker-numbers/.

with any harmful attack. We can see, however, that the proliferation of botnets is a worrisome trend and an important threat to cybersecurity.

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, and then in Section 4 we will go through 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 bought 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 if an outage occurs. Banks encourage consumers and businesses to bank online because the bank experiences massive savings in branch operating costs, even if the interface isn't secure and is regularly exploited by attackers. As pointed out by Anderson and Moore (2006), 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.

There is a natural tension 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 doesn't 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 hard.

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 all banking offline would outweigh the total losses 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 for 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 U.S. Congress on March 20, 2009, AT&T's Chief Security Officer Edward Amoroso estimated that cyber-criminals' annual profit exceeds $1 trillion.[15] That's right, *$1 trillion*. $1 trillion is a lot of money; it's bigger than the entire IT industry, and approximately 7% of U.S. 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 don't 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 most banks aren't revealing how much they're losing to online fraud.[16]

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

However, such a combination of secrecy and FUD (fear, uncertainty and doubt) is dangerous. To understand why, let me first explain how the used car market works. George Akerlof (1970) won a Nobel prize for describing how markets with *asymmetric information*, such as the market for used cars, can fail. Suppose a town has 50 good used cars (worth $2000 each) for sale, along with 50 "lemons" (worth $1000 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 $1500, but at that price no one with a good car will sell, and so the market price quickly ends up near $1000. Consequently, the market is flooded with lemons, since no one with a good car would agree to sell at that price. 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.

Ross Anderson pointed out in 2001 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. A similar effect is triggered by refusing 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 discussion of any 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 we are 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 don't address the lack of reliable information soon, we are liable to end up with decision makers in industry and government refusing to take necessary protections because data explaining the magnitude and nature of the most significant threats just isn't there.

[15]http://commerce.senate.gov/public/?a=Files.Serve&File_id=e8d018c6-bf5f-4ea6-9ecc-a990c4b954c4.

[16]UK banks do report aggregated fraud losses. In 2009, the total reported losses due to all forms of payment fraud were £440 million (approximately $641 million). Of that total, £59.7 million ($87 million) was attributed to online banking losses. Source: http://www.paymentsnews.com/2010/03/uk-card-and-banking-fraud-losses-down-28-in-2009-to-4403mm.html. David Nelson at FDIC has been trying to collect similar figures from U.S. banks on a voluntary basis. He estimates that $120 million was collectively lost by U.S. banks due to malware infections targeting online banking services. Source: http://www.computerworld.com/s/article/9167598/FDIC_Hackers_took_more_than_120M_in_three_months?source=rss_news. In sum, a more accurate estimate of the annual proceeds from online crime is in the neighborhood of the low billions of dollars.

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, but also the platform 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. A secure operating system is more difficult to develop applications for, 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 readily 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 receive widespread adoption. The security benefits of such protocols aren't 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 adopting firms with internal benefits immediately.

Insecurity creates negative externalities. A compromised computer that has been recruited to 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 externalities relevant to cybersecurity is interdependent security. Kunreuther and Heal (2003) 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 (2004) pointed out that free-riding is likely whenever security depends on the weakest link in the chain: firms don't bother investing in security when they know that other players won't 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 next review several different approaches, assessing their suitability to the cybersecurity problem, followed by 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.[17] 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 caused by a failure 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 to these security failures stems from the proliferation of mandatory breach disclosure regulations adopted by many U.S. states (discussed in the next section).

Software companies have long avoided any ex post liability for vulnerabilities in their own products (Barnes 2004). Many have argued that making Microsoft liable for the consequences of exploits targeting Windows would give it a strong incentive to secure it. This is undoubtedly true, but the question is whether it is too blunt an instrument to incent 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 (though 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 toward software liability will also damage the now-flourishing free software community. Graduate students might hesitate to contribute code to a Linux project if they had to worry about being sued years later if a bug they introduced led to a critical vulnerability. Resistance to software liability is one of the few points of agreement between open- and closed-source advocates. Second, 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 don't find them all 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

[17]http://www.rsa.com/products/DLP/ar/10844_5415_The_Value_of_Corporate_Secrets.pdf.

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 (1984) and Kolstad, Ulen and Johnson (1990) 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 harms or is uncertain what minimum standards should be. Unfortunately, both of 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. Facing such a grim reality, we next turn to an alternative approach: information disclosure.

4.1.2 Information Disclosure

Given that information asymmetries are a fundamental barrier to improving cybersecurity, adopting policies that improve information disclosure may be attractive. 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 act. 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 amount and type 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 (Konar and Cohen 1997). The TRI is now available to the public online,[18] and citizens can search the database by ZIP code to learn about chemicals (and the companies which 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 the state of California in 2002.[19] 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.[20] 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 (Mulligan and Bamberger 2007).

Researchers have also found evidence that the information disclosure requirement has both punished violators and reduced harm. Acquisti, Friedman, and Telang (2006) found a statistically significant negative impact on stock prices following a reported breach. Meanwhile, Romanosky, Telang, and Acquisti (2008) 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 an information asymmetry among firms about the prevalence and severity of leakages of personal information. Unfortunately, there is currently no central clearinghouse

[18]http://www.epa.gov/tri/.
[19]California Civil Code 1798.82.
[20]http://www.businesswire.com/portal/site/home/permalink/?ndmViewId=news_view&newsId=20071130005355.

for breach reports, similar to 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.

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 resulting suspicious activity. While consumers may also expect to know about cybersecurity incidents, it is often firms that lack the requisite information on cyber incidents necessary to invest in countermeasures. If the remote login to a power station's controls is compromised and the utility keeps mum about what happened, then other power companies won't fully appreciate the likelihood of attack. When banks don't disclose that several business customers have quickly lost millions of dollars due to the compromise of the company's online banking credentials, the business customers that have not yet fallen victim remain ignorant to the need to take precautions. Thus, in cybersecurity, we face information asymmetries across firms, not only between consumers and firms.

So might information sharing and analysis centers (ISACs) be a viable solution to the asymmetry 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[21] as a way for the federal government to coordinate the protection of critical infrastructures (telecommunications, transport, water, chemical plants, banks, etc.) primarily owned by private industry.

While ISACs have been useful, they are no substitute for a policy of transparency and information disclosure. Many are classified, so any incidents being discussed are kept hidden from those not participating in the meetings, 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 is true, it does underscore the value of the mandatory nature of existing information-disclosure efforts described above.[22] A greater awareness to incidents, even those industries would rather keep hidden, is made possible by mandatory disclosure. Furthermore, in cybersecurity, 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 (Moore and Clayton 2008b).

To wrap up, information disclosure can be a powerful tool in reducing information asymmetries and correcting for misaligned incentives. However, simply righting an information asymmetry won't 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.[23] A robust market for cyber-insurance would offer several key benefits to society. Foremost, insurance could offer a strong incentive to individuals and organizations to take appropriate precautions. Insurance companies could reward security investment by lowering 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

[21]http://www.justice.gov/criminal/cybercrime/white_pr.htm.

[22]Occasionally, particularly egregious incidents are publicized, due to government prodding. For instance, in August 2009 the Financial Services ISAC issued a joint report with the FBI 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. Public disclosures remain the exception, however.

[23]See Boehme and Schwarz (2010) for a complete account of cyber-insurance's prospects and limitations.

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 (Bandyopadhyay, Mookerjee, and Rao, 2009). Furthermore, the biggest benefits ascribed to cyber-insurance have not been realized. Rather than differentiate premiums by observed security levels, insurance companies base premiums on non-technical criteria such as firm size. Additionally, insurance companies have not amassed a large claims history documenting incidents.

Why has the market for cyber-insurance been such a disappointment? Factors on both the demand and supply side offer explanation. On the demand side, insurers complain of a lack of awareness to cyber-risks by firms. In fact, they point to mandatory breach disclosure legislation (as described in the previous section) as a significant step in the right direction, arguing that it has increased awareness at the executive level of this one particular category of threat. Consequently, policies that increase disclosure of cyber risks and incidents would help stimulate further growth in the cyber-insurance market. However, not all demand-side challenges can be dealt with by increased awareness alone. 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 PC infections when they are not on the hook for mitigation. Legislation that clarifies liability for cyber incidents would go a long way toward 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—help explain why insurance companies still don't differentiate premiums based on technical criteria. Certification schemes might help, but designing security certifications that cannot be gamed is hard. Examples of failed certifications include Common Criteria-certified "tamper-proof" PIN entry devices broken by cleverly-placed paper clips (Drimer, Murdoch and Anderson 2008) and more malicious websites receiving the TrustE seal of approval than legitimate sites (Edelman 2009). The other big supply-side problem is that losses from many types of information security risks are globally correlated. Given Windows' dominant market share, a new exploit that compromises Windows PCs 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 (Boehme and Kataria 2006). 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 its employees. Why would indirect liability ever be desirable? Following the logic of Lichtman and Posner (2004), 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 couldn't pay up even if caught. Second, high transaction costs could make designing contracts that dish out 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 defamatory content contributed by its users. Until the CDA was passed, service providers were reticent to moderate any posts from users out of fear that doing so would expose them to liability for all content contributed by users. Section 230 of the CDA offered immunity to service providers that chose to voluntarily delete contributions from users 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 users that violate copyright online. The DMCA also exempts service providers from liability for copyright infringement carried out by its customers. However, this time there's 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 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 put in place procedures to stop Internet gambling transactions. Because all Internet gambling operations needed credit card payments to make money, 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 a range of other online ills, including child pornography, controlled substances and tobacco sales to minors. MacCarthy (2009) offers 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 operating 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 (MacCarthy 2010). For frauds occurring at brick-and-mortar stores, banks traditionally foot the bill, not the merchants where the fraud occurred. 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 at 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 systems. A key take-home point when considering what to do about cybersecurity more broadly is that legal ambiguity about which intermediary must pay for remedies is undesirable and can lead to nasty legal battles.

To sum up, Congress has acted to regulate the undesirable activities of online users by articulating what intermediaries can or must do. There's 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 Recommendation 1: Mitigating Malware Infections via ISPs by Subsidized Cleanup

As described in Section 2.4, botnets comprising computers 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. Hence, 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 its customers are infected with malware by helping to coordinate the cleanup of affected computers. In exchange for cooperation, ISPs receive exemption from liability for the harm caused by the infected machines. If ISPs do not cooperate, then ISPs become liable for the harm caused by infected machines.
- The costs of cleanup will be shared between ISPs, government, software companies and consumers.
- Reports of infections (including ISP, machine OS type, infection vector, time to remediation, remediation technique) must be reported to 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 affecting its software. Software companies receive 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 dollar amount. Consumers are guaranteed that they will not be disconnected by their ISP 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 place 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.

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 doesn't work because most malware tries to disable new updates to the anti-virus software. Another option for Windows users comes through Windows Update. While far from complete, Microsoft's Malicious Software Removal Tool (MSRT) does automatically detect and remove popular types of malware. If these precautions fail, then the user often remains completely ignorant of the malware's presence. However, most malware-infected computers leave a trail of malicious activity that can be identified by third-party security companies which monitor Internet traffic. These companies often notify the relevant ISP of the activity. Some ISPs also actively detect computers that participate in botnets.[24] They then pass along 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.

[24]http://www.maawg.org/sites/maawg/files/news/MAAWG_Bot_Mitigation_BP_2009-07.pdf.

Once notified of malware on their customers' computers, ISPs have several options for taking action. At a bare minimum they can pass along the notice to consumers. In October 2009, Comcast announced a trial program to notify customers that they are infected via a browser pop-up, with links to instructions for removal.[25] Such notification-only schemes rely on customers to take the necessary steps, which sometimes works for tech-savvy users and on types of malware detectable by tools such as Microsoft's MSRT. Inevitably, though, malware is often not removed by users after they have been notified. For these cases, Comcast has partnered with McAfee to offer a remediation service by a skilled technician for $89.95. Australian ISPs recently announced a notification-based effort for all its ISPs.[26]

Another ISP-based option is to place infected computers into "quarantine." Once in quarantine, users are required to download and install anti-virus software and malware removal tools. They are then only permitted to rejoin the wider Internet once 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 the ISP and more customer-support calls are made. Some ISPs use quarantine systems, but even those that do only use them for a minority of affected customers. Recently Dutch ISPs announced a signed agreement to notify and quarantine affected customers.[27] Note that in both the Dutch and the 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 to notification that customers are infected with malware is to take no action. Why? The incentive for ISPs to intervene is very weak (van Eeten and Bauer 2008). Malware harms many victims, from consumers whose credentials are stolen to the targets of DDoS attacks. However, ISPs are not affected very much, apart from the prospects of being chided by other ISPs if too many customer machines are sending 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 that come in after sending out notices or placing customers into quarantine are very high. For the ISP, it is far less costly to ignore the notifications.

Consequently, the status quo of malware remediation is unacceptable. Many ISPs choose not to act, and even those that do avoid cleaning up the hard cases. Notification-only approaches leave many computers infected, while quarantine-based schemes can unfairly shut off the Internet connections of consumers that 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 actually carrying out 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 (2004) argue for ISPs to take on strict liability for the actions of its customers' computers. In other words, they suggest simply making the ISPs take the blame for malware-infected customers, and let them choose how they remedy the situation given the threat of legal responsibility. Given the history of exemptions ISPs have secured from responsibility for the actions of its customers in other contexts, we find such an aggressive approach 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 actively intervening voluntarily, clarifying the legal right to do so might embolden wary ISPs to act. However, there are distinct

[25]http://www.comcast.com/About/PressRelease/PressReleaseDetail.ashx?prid=926.
[26]http://iia.net.au/index.php/section-blog/90-esecurity-code-for-isps/757-esecurity-code-to-protect-australians-online.html.
[27]http://www.darkreading.com/blog/archives/2009/09/dutch_isps_sign.html.

disadvantages of this approach. Notably, it does nothing to compensate for the weak incentives 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 notification first, and continue ratcheting up efforts if notification fails to fix the problem. Quarantine may be tried next, followed by perhaps sending a technician to remediate the machine. Any legislation wouldn't be prescriptive in laying out the steps that must be tried and their order; rather, the scheme should be flexible in allowing ISPs to try different approaches so 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 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. (2008) recommended that the European Commission introduce fixed penalties for ISPs that do not expeditiously comply with notifications of compromised machines present on 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, 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 its infected customers are cleaned up will impose a costly obligation on the ISP. This is somewhat unfair, since it is not the ISP's fault that the user has been infected. Yet indirect liability regimes need not be fair to be effective. However, a fair allocation of responsibilities is helpful to 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 it created. To that end, we recommend that the costs of cleanup should be shared between ISPs, government, software companies and consumers. ISPs already pay by 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 is already precedent for cost-sharing between third parties in the cybersecurity context. First, Luxembourg is exploring the possibility of subsidizing malware cleanup (Clayton 2010). 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 its breach affecting 46 million cardholders,[28] while in January 2010 Heartland agreed to pay MasterCard $41 million following its breach of 100 million credit card numbers.[29] Rather than negotiating one-off settlements between intermediaries, we recommend establishing a fund to receive regular payment 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 par-

[28]http://www.businesswire.com/portal/site/home/permalink/?ndmViewId=news_view&newsId=20071130005355.
[29]http://www.pcworld.com/businesscenter/article/196711/heartland_mastercard_settle_over_data_breach.html.

ticipation more palatable, we recommend that by helping to pay for the cleanup software companies be granted safe harbor from any harm the compromised machines have caused prior to cleanup. How much should companies pay? Payment could be distributed according to what caused the infections. If infection reports included the method of exploitation, then it is easy to figure out whether the culprit is Windows XP (Microsoft pays) or Acrobat (Adobe pays). Once the scheme is up and running, contribution amounts for the next quarter can be based upon 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, it will be difficult to track down the party responsible for developing the software that has been exploited (e.g., if the software is open source). 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 compromise and the cleanup process. Reports of infections (including ISP, machine OS type, 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,[30] for instance.

Mandatory collection and publication of data is an essential component of the scheme and part of the grand bargain between ISPs and software companies receiving liability exemptions in exchange for cooperation with the cleanup process. It's not there just to help researchers. Mandatory disclosure of infections will help fix the information asymmetry plaguing information security investment (described in Section 3.2). Disclosure will put valuable security incident data in the public domain, and it is likely that it will trigger a similar "sunshine effect" as has been observed in environmental pollution due to the Toxic Release Index and in protecting personal information due to breach-disclosure laws. Some of the worst offenders (both ISPs and software companies) will be uncovered, raising awareness to the problem and providing an incentive for investment in defense. Progress will become measurable, not only to insiders but also to outside perspectives on the comprehensiveness 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 a government level. If the U.S. 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 of the best practices may become infected. According to Panda Security, 3.94% of U.S. 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 very 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 furthermore that 1.3% of Google's incoming search queries return at least one drive-by-download link in its results (Provos et al. 2008).

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 fix the problem. Consumers are generally not in a good position to defend themselves. They don't write the buggy software, and so they can't plug the

[30]http://xml.coverpages.org/iodef.html.

holes; they don't have a network-level view of Internet traffic, so they can't determine whether they are infected (as ISPs can). At best, they can take some safety precautions such as patch their computers and install antivirus. There's little more we can expect from them, and even if we got all consumers to automatically install patches and run antivirus software, we'd still have a problem. Consequently, consumers are not in the best position to cheaply fix the problem.

In light of this reality, any resulting policy should focus on ensuring that consumers are protected in the course of any cleanup efforts. Consequently, we recommend that any 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 at the same time minimizing any moral hazard among some users. Perhaps the fee could be slightly higher for users that do not have antivirus software installed. It is also essential that the burden on the ISP is to actually remedy the infection. Disconnecting users' Internet connections is not an acceptable remediation, given the increasing reliance on the Internet to provide basic services. The only exception allowing disconnection could be if consumers do not cooperate with the cleanup efforts of the ISP. Otherwise, ISPs should have a duty to cleanup, capping the out-of-pocket expenses for consumers. This is to address concern that ISPs will choose to kick off users rather than clean them up (in Lichtman and Posner's words, "concern that liable ISPs will be overly cautious and thus inefficiently exclude marginal subscribers").

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 to regularly publish the following aggregated loss figures related to online banking and payment cards on data.gov:

- Incident figures: # of incidents, total $ stolen, total $ recovered for specified # of incidents
- Victim bank demographics: # banks affected, # customer accounts impacted per bank, $ lost per customer, bank type, precautions taken by bank (2-factor authentication, back-end controls used)
- Victim customer demographics: business v. consumer breakdown—#s and losses
- Attack vector (if known): keyloggers, phishing, card skimming, payment network compromise, etc.
- Business category: online banking, payment cards (transaction type: retail, card present, card not present), ATM fraud

At present, no objective measures exist to answer 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? Is 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: ask financial institutions to report back on fraud losses and aggregate their responses. It is not as though such 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.[31] Richard Sullivan, economist at the Federal Reserve, has argued that fraud statistics should be published in order to get a

[31]http://www.paymentsnews.com/2010/03/uk-card-and-banking-fraud-losses-down-28-in-2009-to-4403mm.html.

better grip on fraud levels and inform whether investments to secure the payment card infrastructure are needed (Sullivan 2009).

Within the U.S., there are some existing efforts to collect data on online frauds. David Nelson at the FDIC has been trying to collect fraud figures from U.S. banks on a voluntary basis. He estimates that $120 million was collectively lost by U.S. banks due to malware infections targeting online banking services.[32] The FBI runs the Internet Crime Complaint Center (IC3), which invites members of the public to submit reports of a wide variety of Internet scams. A few aggregate figures from the IC3 report are regularly made available in annual reports,[33] but most access to the IC3 data is restricted 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 we recommend. First, the reports are voluntary in nature, making them incomplete, unrepresentative, and impossible to draw reliable trends from. Very few privacy breaches were disclosed until the California law was passed, and we might suspect that the reports of online fraud are also inaccurately represented. In the case of IC3, the trouble is that quantifying losses is difficult for many circumstances, as this too relies on self-reporting. Second, they are often secret in nature—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* 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 and Z, but businesses are not, and merchants are expected to share responsibility for covering 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 who pays what share. 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 both consumers and businesses. Additionally, we should learn 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, whether any groups pay a disproportionate share, and whether 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 existing user base is enormous, and presents a significant barrier to would-be new entrants offering a more secure payment alternative. Having already invested heavily in a less secure payment technology and achieved market dominance, existing payment networks may be reticent to invest further in security mechanisms to reduce fraud that is borne in part by third parties. Payment networks might reasonably retort that they are investing in security, and point to efforts already undertaken in Europe to upgrade to PIN-based smartcard authentication.

A credible reporting of financial fraud losses can settle any dispute 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 (MacCarthy 2010). Comprehensive fraud statistics would help answer to banks *and* merchants whether there has been a substantial enough increase in card-not-present fraud to justify further security investment. Similarly, the National Security for Secure Online Transactions[34] being pitched by the White House needs buy-in from the private sector

[32]http://www.computerworld.com/s/article/9167598/FDIC_Hackers_took_more_than_120M_in_three_months?source=rss_news.

[33]http://www.ic3.gov/media/annualreport/2009_IC3Report.pdf.

[34]http://pindebit.blogspot.com/2010/04/national-strategy-for-secure-online.html.

to be successful. To get that buy-in, firms need to believe that improved online authentication is needed. How can firms agree to spend on security when they do not have an accurate picture as to how much is being lost due to the less secure infrastructure we have today? Publishing regular statistics on losses now will motivate future investment if the problem is truly as big as has been claimed.

4.4 Recommendation 3: Mandated Disclosure of Control System Incidents and Intrusions

We have received stark warnings from anonymous intelligence officials that the Chinese and Russians have regularly intruded into the U.S. electrical grid.[35] Yet no documented case of a successful cyber attack on process control systems has been publicly presented. In fact, when researchers from the Tuck School of Business interviewed an oil and gas refiner as part of a field study (Dynes, Goetz and Freeman 2007), 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 only consider investing in process control systems security after a similar-sized refinery was attacked first.

Such different perspectives are hard to rectify—that attacks are already pervasive yet operators on the ground have yet to observe a single incident. One possible explanation is that the reports of incidents are exaggerated. Many of those sounding the alarm do certainly stand to gain from increased security investment. 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 the industry, or there hasn't been much to report.

Fortunately, there is a reasonable way to get to the bottom of this conundrum: adopt mandatory disclosure of all cyber incidents and intrusions, with a substantial public reporting capacity. If the intrusions are in fact happening, those who detect them should have a duty to report these. In fact, the ISACs could serve as the organization to receive reports, provided that there is a clear duty to produce public reports that receive widespread dissemination.

Recommendation 3: Mandatory disclosure of control system incidents and intrusions to the relevant ISACs, who provide further public dissemination.

There has been some tentative movement in this direction within the electricity industry. The self-regulatory body NERC has required power companies to start reporting to regulators any time 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, not shared with other firms in the industry. This is a useful start, as 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 into the SCADA system need not be reported, unless it caused the power to go out. It also has no explicit requirement to share the reported information with other utilities, which doesn't solve the problem of the oil refiner who is waiting to invest until he hears about others 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-tail nature of cyber attacks on process control systems means that the effort could yield few reports. Furthermore, the problem of externalities remains.

[35] http://online.wsj.com/article/SB123914805204099085.html.

4.5 Recommendation 4: Aggregate Reports of Cyber Espionage and Report to WTO

Industrial espionage is claimed to be a significant problem for U.S. companies. However, these companies are naturally reticent to publicly discuss their experiences of espionage out of fear that their stock price may take a hit. Perhaps, though, the thinking is starting to change. In January 2010, Google disclosed that they had been the victim of a cyber attack apparently originating in China whose purpose was industrial espionage.[36] Subsequently it was revealed that at least 34 companies were affected, including Yahoo, Symantec, Northrop Grunman and Dow Chemical.[37]

Unfortunately, since the trade secrets were believed to be stolen by someone internationally, the Uniform Trade Secrets Act and Economic Espionage Act cannot easily be enforced. This does, however, leave one option: the TRIPS agreement of the World Trade Organization. Deciding to bring cases to the WTO is always politically delicate. However, if the U.S. suspects that industrial espionage is rife, and largely coming from a single country (i.e., China), then it may be worth preparing a complaint to the WTO. It's true that such a complaint could potentially harm the stock prices of the firms named victims. 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 U.S. to take action.

5 CONCLUSION

An economic perspective is essential for understanding 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 focus on getting Internet service providers 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 to cybersecurity issues and assign responsibility for action within the private sector so that the risks to society may be mitigated.

REFERENCES

Acquisti, A., A. Friedman, and R. Telang. 2006. Is There a Cost to Privacy Breaches? An Event Study. *Proceedings of the International Conference on Information Systems (ICIS)*, Milwaukee, WI.

Akerlof, G. A. 1970. The Market for 'Lemons': Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics* 84(3):488-500.

Anderson, R. 2001. Why Information Security is Hard—An Economic Perspective. *Proceedings of the 17th Annual Computer Security Applications Conference*, pp. 358-65. IEEE Computer Society.

Anderson, R., and T. Moore. 2006. The Economics of Information Security. *Science* 314(5799):610-13.

Anderson, R., R. Boehme, R. Clayton, and T. Moore. 2008. Security Economics and the Internal Market. European Network and Information Security Agency. Available at http://www.enisa.europa.eu/act/sr/reports/econ-sec/economics-sec/at_download/fullReport.

Bandyopadhyay, T., V. Mookerjee, and R. Rao. 2009. Why IT managers don't go for cyber-insurance products. *Communications of the ACM* 52(11):68-73.

Barnes, Douglas A. 2004. Deworming the Internet. *Texas Law Review* 83(1), available at http://ssrn.com/abstract=622364.

Boehme, R. and G. Kataria. 2006. Models and measures for correlation in cyber-insurance. *Proceedings of the Workshop on the Economics of Information Security*, University of Cambridge, UK.

Boehme, R. and G. Schwarz. 2010. Modeling Cyber-Insurance: Towards a Unifying Framework. *Proceedings of the 9th Workshop on the Economics of Information Security*, Cambridge, MA, June 7-8. Available at http://weis2010.econinfosec.org/papers/session5/weis2010_boehme.pdf.

Camp, L. J., and C. D. Wolfram. 2004. Pricing Security: A Market in Vulnerabilities. In Economics of Information Security, Vol. 12, *Advances in Information Security*, ed. L. J. Camp and S. Lewis, 17-34. Boston: Kluwer Academic Publishers.

[36]http://www.darkreading.com/vulnerability_management/security/attacks/showArticle.jhtml?articleID=222700786.

[37]http://www.washingtonpost.com/wp-dyn/content/article/2010/01/13/AR2010011300359.html.

Clayton, R. 2010. Might Governments Clean-up Malware? *Proceedings of the 9th Workshop on the Economics of Information Security*, Cambridge, MA, June 7-8. Available at http://weis2010.econinfosec.org/papers/session4/weis2010_clayton.pdf.

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

Dynes, S., E. Goetz and M. Freeman. 2007. Cybersecurity: Are Economic Incentives Adequate? In *Critical Infrastructure Protection*, pp. 15-27. Springer.

Edelman, B. 2009. Adverse selection in online "trust" certifications. *Proceedings of the 11th International Conference on Electronic Commerce*, pp. 205-212. ACM Press.

Information War Monitor. 2009. Tracking GhostNet: Investigating a Cyber Espionage Network. Available at http://www.scribd.com/doc/13731776/Tracking-GhostNet-Investigating-a-Cyber-Espionage-Network.

Kolstad, C., Ulen, T. and Johnson, G. 1990. Ex Post Liability for Harm vs. Ex Ante Safety Regulation: Substitutes or Complements? *American Economic Review* 80(4):888-901.

Konar, S., and M. Cohen. 1997. Information As Regulation: The Effect of Community Right to Know Laws on Toxic Emissions. *Journal of Environmental Economics and Management* 32(1):109-124.

Lichtman, D.G. and E.A. Posner. 2004. Holding Internet Service Providers Accountable. In *The law and economics of cybersecurity*, ed. Mark F. Grady, F. Paris, pp. 221-258.

MacCarthy, M. 2009. What Internet Intermediaries Are Doing About Liability and Why It Matters. *ExpressO*. Available at http://works.bepress.com/mark_maccarthy/1.

MacCarthy, M. 2010. Information Security Policy in the U.S. Retail Payments Industry. 9th Workshop on the Economics of Information Security, Cambridge, MA, June 7-8. Available at http://weis2010.econinfosec.org/papers/panel/weis2010_maccarthy.pdf.

Moore, T., and R. Clayton. 2007. Examining the Impact of Website Take-down on Phishing. *Proceedings of the Anti-Phishing Working Group eCrime Researchers Summit*, pp. 1-13.

Moore, T., and R. Clayton. 2008a. The Impact of Incentives on Notice and Take-down. In *Managing Information Risk and the Economics of Security*, ed. M. Eric Johnson, 199-223. New York: Springer.

Moore, T., and R. Clayton. 2008b. The Consequence of Non-cooperation in the Fight against Phishing. *Proceedings of the Anti-Phishing Working Group eCrime Researchers Summit*, pp. 1-14.

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

Nagaraja, S., and R. Anderson. 2009. The snooping dragon: social-malware surveillance of the Tibetan movement. *University of Cambridge Computer Laboratory Technical Report UCAM-CL-TR-746*. Available at http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-746.pdf.

Provos, N., P. Mavrommatis, M. Rajab, and F. Monrose. 2008. All Your iFrames Point to Us. *Proceedings of the USENIX Security Symposium*, pp. 1-15.

Romanosky, S., R. Telang, and A. Acquisti. 2008. Do Data Breach Disclosure Laws Reduce Identity Theft? *7th Workshop on the Economics of Information Security*, Hanover, NH. Available at http://ssrn.com/paper=1268926.

Shavell, S. 1984. A Model of the Optimal Use of Liability and Safety Regulation. *RAND Journal of Economics* 15(2):271-280.

Stern, H. 2009. The Rise and Fall of Reactor Mailer. *Proceedings of the MIT Spam Conference*, Cambridge, MA, March. Available at http://projects.csail.mit.edu/spamconf/SC2009/Henry_Stern/.

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

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